UNIVERSITY OF CALIFORNIA, SAN DIEGO

Statistical Discrimination, Recognition and Altruism, and Pure/Mixed Strategy Manipulation

A dissertation submitted in partial satisfaction of the requirements for the degree

of Doctor of Philosophy

in

Economics

by

Bryan Tomlin

Committee in Charge:

Professor Gordon Dahl, Chair

Professor Vincent Crawford, Co-Chair

Professor Kate Antonovics

Professor Uri Gneezy

Professor Craig McKenzie

2010

The Dissertation of Bryan Tomlin is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

Co-Chair

_____

Chair

University of California, San Diego

2010

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Chapter 2

Chapter 3

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Gordon Dahl, for his guidance and support as my advisor and as the chair of my committee. Without his help, this dissertation would not exist. I would also like to thank Professor Vincent Crawford and Professor Kate Antonovics for their advice and wisdom over the course of my graduate studies. Finally, I would like to thank my wife Sarah for absolutely everything, she is my world.

Chapter 1: I would also like to acknowledge Michael Ewens and Choon Wang, who co-authored with me the first chapter of this dissertation. Their assistance was invaluable in the authoring of this chapter.

Chapter 2: I would also like to acknowledge the Russell Sage Foundation for their support which made this research possible. Without their small grant for behavioral economics, there is no way that I would have been able to conduct this experiment.

VITA

DOCTORAL STUDIES:
PhD University of California, San Diego, 2004 to 2010

UNDERGRADUATE STUDIES:
Bachelor of Arts, Economics, University of California, Berkeley, with honors, 2004
Bachelor of Arts, Political Science, University of California, Berkeley, 2004

TEACHING EXPERIENCE:
Austin Community College - Instructor
Microeconomics - Summer 2009

Park University - Instructor
Statistics for Business and Economics - Summer 2009

Rady School of Management at University of California, San Diego - Teaching Assistant
Courses: Strategy

University of California, San Diego - Teaching Assistant
Courses: Microeconomics, Intro/Intermediate/Advanced Econometrics, Game Theory, Choice Under Uncertainty, Financial Economics, Labor Economics, Economics of Industrialization, Economic History of the United States, Environmental Economics

PROFESSIONAL EXPERIENCE:
M. Davis and Company Inc. – Economic Consultant
January 2010 - Present

HONORS, SCHOLARSHIPS, AND FELLOWSHIPS:
Russell Sage Foundation - Small Grant for Behavioral Economics
Awarded Fall 2008

Institute for Applied Economics – Research Grant
Awarded Fall 2008

UC San Diego *– Department of Economics Teaching Assistant Excellence Award*
Awarded Fall 2009

UC San Diego *- Department of Economics Heller Prize* – Outstanding 3rd Year Paper
Awarded Spring 2007

B.P. Foundation Scholarship
Awarded 2004, 2005, 2006, 2007, 2009

RESEARCH PAPERS:
Statistical Discrimination in the U.S. Apartment Rental Market

The Importance of Appearing Earnest: Altruism in the Presence of Misunderstandings

Choosing Between Mixed and Pure Strategies in a Simple Experiment

PROFESSIONAL ACTIVITIES:
Journal Referee: *American Economic Review*

ABSTRACT OF THE DISSERTATION

Statistical Discrimination, Recognition and Altruism, and Pure/Mixed Strategy Manipulation

by

Bryan Tomlin

Doctor of Philosophy in Economics

University of California, San Diego, 2010

Professor Gordon Dahl, Chair
Professor Vincent Crawford, Co-Chair

This dissertation is comprised of three autonomous chapters all of which have one thing in common: they utilize experimental manipulations to answer questions of interest to economists and/or society. Chapter 1 finds significant evidence of discrimination against African Americans in the apartment rental market throughout the entire United States – subtle discrimination which would likely go unnoticed by any individual, though the effects are likely felt by many individuals. Chapter 2 finds evidence of people actively exploiting others' perceptions in order to obtain a favorable appearance – subjects will behave selfishly while attempting to appear selfless or egalitarian. Chapter 3 shows that a simple manipulation in the wording of a problem can result in subjects playing strategies more closely aligned with Nash Equilibrium predicted behavior.

CHAPTER 1

STATISTICAL DISCRIMINATION IN THE U.S. APARTMENT RENTAL MARKET: EVIDENCE FROM A LARGE SAMPLE FIELD STUDY

Michael Ewens	Bryan Tomlin*	Choon Wang
U.C. San Diego	U.C. San Diego	U.C. San Diego
Department of Economics	Department of Economics	Department of Economics

MAY 2010

Abstract

We test for statistical discrimination in the apartment rental market using an email-based correspondence study. Emails containing randomly assigned white or black sounding names are sent to over 14,000 vacancy listings posted on Craigslist.org across 35 U.S. cities. By manipulating the level of positive or negative information included in inquiry emails sent to these postings we are able to evaluate the effect of information on racial discrimination. We find evidence of preferential treatment of whites in terms of the likelihood of receiving a response in a baseline treatment where no additional information is provided. The racial gap in differential treatment is unchanged when positive information is added, though the gap diminishes in the presence of negative information.

I. Introduction

Over the past decade the number of discriminatory complaints filed with both public and private fair housing organizations has not diminished. The U.S. Department of Housing and Urban Development (HUD), the U.S. Department of Justice (DOJ), state and local government organizations, and private fair housing groups receive nearly 30,000 complaints each year.[1] This number is thought to represent just a small percentage of the total number of annual discriminatory acts in the housing market as many go unreported. HUD conservatively estimates the number of housing discrimination incidents to be approximately two-million per year.[2] Some claim that such behavior is not the result of prejudice but rather of profit maximization, and landlords are simply using an applicant's race as an indicator of their expected quality as a tenant. This explanation of differential market outcomes by race is referred to as "statistical discrimination." According to this theory, were the landlord to be presented with useful information about an applicant, such as their vocation, race would become relatively less informative and differential treatment would decrease relative to when the landlord knew nothing about the applicants. Data limitations make empirical tests of these claims difficult. Researchers can observe where individuals choose to live, but not the choice set they were presented with in making this choice, nor the information they presented to landlords when applying for apartments. HUD attempts to address this issue using audit studies which send out actors as prospective tenants/home-buyers in search of a home. Such studies, however, are limited by surveyor bias (auditors seeking out or misreporting discrimination to support the researcher's agenda or their own political view) and the fact that – no-matter how well trained – paired auditors are likely to be heterogeneous in ways which may be insignificant or

[1] United States Department of Justice, Civil Rights Division (www.justice.gov/crt)
[2] www.HUD.gov

unobservable to researchers, but important and easily observable to landlords and real-estate agents.

To circumvent these data limitations and study discrimination in the apartment rental market we conduct a field experiment using a detailed correspondence methodology.[3] We send out over 14,000 inquiry emails to apartment listings posted on craigslist.org in 35 cities across the continental United States over the period of two months. We observe all information presented to applicants by landlords while experimentally manipulating information presented to landlords by applicants. Specifically, we control any and all indicators of the applicant's race, gender, and potential quality as a tenant, as well as the exact words they say to landlords. This ensures that our applicants are perfectly homogenous and any difference in treatment must be the result of one of our experimentally altered information/race controls. Most notably, since we are able to manipulate the quality of each applicant, we are able to cleanly examine whether discrimination in the apartment rental market is possibly the result of statistical discrimination.

The race of the inquirer (applicant) is altered by randomly assigning to each inquiry a fictitious black or white-sounding applicant name and corresponding email address. A male or female sounding name is randomly assigned to each inquiry email with equal probability.

The quality of each applicant is manipulated by randomly inserting a statement disclosing positive or negative information into an applicant's inquiry. "Positive information" takes the form of a statement indicating that the applicant holds a steady job in a respected vocation and is a non-smoker. "Negative information" is presented by stating that the applicant has a bad credit rating and smokes. Approximately 30% of the inquiries serve as a baseline treatment in which the applicant simply states interest in the unit and asks about whether or

[3] Correspondence studies are conducted entirely through writing. Our methodology utilizes emails between our fictitious applicants and real landlords who posted rental listings on craigslist.org as the medium of correspondence.

not the listed unit is available. Another 40% do the same (using the exact same scripts) but also include the positive information statement. The final 30%, again using the same baseline scripts, also include the negative information statement.

We measure the effects of our treatments in terms of two landlord interest metrics: the likelihood of receiving a response, and the likelihood of receiving a positive response. We confirm previous findings of anti-black discrimination in the housing market in terms of both metrics. In the baseline condition (no additional information) black applicants receive approximately 9 responses for every 10 received by white applicants, and 8.4 *positive* responses for every 10 received by white applicants. This differential treatment is highly significant ($p<0.001$).

We find that the addition of positive information to inquiries does nothing to diminish the preferential treatment of white applicants relative to black applicants: the ratio of black to white positive responses becomes 8.5 to 10 in the presence of positive information. This result is surprising since it indicates that African American professionals are treated just as negatively by landlords as African Americans about whom landlords know nothing. Negative information is found to significantly decrease the absolute difference in positive response rates between black and white applicants, though the relative difference in treatment (as a ratio) is actually increased: the black to white positive response rate ratio become 8 to 10 in the presence of negative information.

Section II discusses previous research concerning racial discrimination in both the housing and employment markets. Section III presents our experimental methodology. Section IV presents and interprets the results, and section V concludes.

II. Review of Relevant Literature

Gathering data on discrimination in the housing market is difficult since the qualities of an ideal tenant are difficult to measure. In the interest of controlling for these hard to measure qualities, audit studies conducted by HUD train auditors (actors sent out to apply for housing in the same neighborhoods) in pairs to behave exactly like one another, ideally leaving ethnicity as their sole difference. This training includes preparation about fictitious back-stories, resumes, and demeanor. Devah Pager (2007) offers insight, summary, and analysis of the ways in which audit studies are conducted. Heckman and Siegleman (1993) and Heckman (1998) offer compelling arguments against the use of field studies, noting that auditors may be inclined to alter their behavior so as to illicit the sort of results they believe researchers are looking for. They also suggest that auditors, regardless of their training, may differ in important ways, unobservable to researchers but obvious to landlords and real-estate-agents.

Yinger (1986) examines discrimination in the real estate market taking the number of housing units shown and/or recommended to paired black and white auditors in the Boston area as his dependent variable. He finds black auditors to be shown or made aware of 30 percent fewer units than white auditors. Page (1995) tests discrimination in a number of cities and between black-white and Hispanic-white pairs using a fixed effects Poisson model, finding that blacks and Hispanics are shown 10 to 20 percent fewer units than are white auditors, and that this discrepancy varies between cities. This model has since been the standard in audit literature since it recognizes that the sequential manner in which auditors visit brokers may have a non-negligible effect on outcomes. Our research avoids this complication by sending only one inquiry to each landlord. The downside of our approach, however, is that we are unable to identify differential treatment by race taking place within any one landlord observation.

Roychoudry and Goodman (1996) examine differences in units shown to auditors in the Detroit area, keeping track of real estate agent characteristics. They confirm, once again, that blacks were shown significantly fewer units, and that this discrimination is less common among black agents and more common among older agents, illustrating that dimensions other than auditor heterogeneity may play an important role in explaining discrimination. This points to one weakness of our study since we are unable to identify specific landlord characteristics.

Ondrich et al (1998) and Ondrich et al (1999) study the probability that an auditor will be shown a unit in various large cities using a discrete choice model, using data from brokers as well as landlords. These studies show that brokers and landlords discriminate due to their own prejudices as well as the prejudice of their white clients/tenants. Ondrich et al. (2003) go on to show that a common form of discrimination against blacks takes the form of "redlining" wherein brokers withhold units in integrated neighborhoods from some home seekers. We also observe a significant amount of differential treatment being manifested by a simple decrease in the response rate to black applicants, indicating that some discrimination may be the result of landlords ignoring inquiries from African Americans.

The name generation process utilized by this paper follows Bertrand and Mullainathan (2005) (BM2005). BM2005 gathered potential first names using the birth certificate data of all babies born in Massachusetts between 1974 and 1979. This allowed them to select names which denoted a certain race with a very high probability. Additionally, BM2005 conducted a survey to ensure that a random person would be likely to select the true ethnicity of a person based solely on their first name. The African-American names generated were quite common amongst the population. This not only ensured that their names covered a large segment of the population, but also that it avoided some of the potential issues with unique-African-American

names as described in Fryer and Levitt (2004). Given the high quality of BM2005's name generation, we borrow from their name pool and utilize the same first names used in their study (see Table C1).

A correspondence study of housing market discrimination which this study will use as a baseline for comparison is that of Carpusor and Loges (2006). These authors sent out the same generic email inquiry to 1,115 landlords advertizing apartment vacancies in the Los Angeles area over a 10 week period, but randomly assigned one of three different male names to each inquiry (black, white, or Arab-sounding). The rent of the unit to which each applicant was applying was recorded, and data was analyzed using a comparison of means. We find similar black/white results in our baseline treatment and build upon these results by altering the quality of applicants to test for statistical discrimination.[4]

III. Methodology

Given the magnitude and level of automation required by this study, the full (detailed) methodology is reserved for the appendix. The methodology presented below is intended to be sufficient for understanding our results.

A. Email Generation

This being a correspondence study, applicant characteristics are conveyed to landlords entirely through email text. As such, the only signal of race and gender is the applicant's name. To maximize the probability that landlords will observe this signal, the applicant's full name is

[4] Another correspondence study of rental market discrimination was carried out by Ahmed and Hammarstedt (2008). Their study, however, takes too many departures from our own to provide a suitable base of comparison. The three most notable departures are as follows: First, they attached the same name to the same inquiry every time, making it impossible to determine whether applicant name or applicant inquiry was being discriminated against. Second, their study was conducted in Sweden, where nearly half of the rental dwellings are owned by the government. And finally, they compared only white male, white female, and Muslim male applicants (finding the white female most likely to receive an invitation to view an apartment (53%) followed by the white male (41%) and finally the Muslim male (18%).

presented three times in every email: first in the email address, which is always of the form "first.last@domain.com," second in the intro sentence of the email text, which always begins, "My name is 'First Last'…," and third in the closing signature of the email. Names were generated using the same pool utilized by BM2005 in their correspondence study, combined with the US Census Name Survey. Names resulting from this combination include: Allison Bauer, Ebony Washington, Matthew Klein, and Darnell Booker.[5]

Below are two samples of emails we sent. These samples, unlike the actual emails we sent, include numbers within parentheses labeling the five key "elements" which composed the texts of our emails. The purpose of each element is explained below.

Example 1, Positive Treatment –

> (1) Hello,
>
> (2) My name is [Full Name], and I am writing in response to your listing for an apartment for [apartment rent]/month. (3) In case you're interested, I do not smoke and I work full time as an architect. (4) Is this unit still available? (5) Thank you for your time,
>
> [Full Name]

Example 2, Negative Treatment –

> (1) Hi,
>
> (2) My name is [Full Name], I am responding to your craigslist posting for an apartment listed at [apartment rent]/month. (3) Just so you know, I am a smoker and my credit rating is below average. (4) I realize places go fast sometimes, is this unit still available? (5) Thank you,
>
> [Full Name]

Each email text was generated by randomly selecting the text for each of the five elements (numerated in the sample emails above). With the exception of the statement of

[5] White female, black female, white male, and black male respectively. For a full list of names used, see Table C1 of the appendix.

quality, all text was pulled from the same pools. (1) is an introductory *hello* statement. (2) is a *statement of interest* in the apartment which always includes the rent of the unit being applied to (to avoid confusion in case the landlord has posted multiple listings). (3) is a *statement of quality* which is randomly included (or not included) to define our treatments. This element is discussed in detail below. (4) is an *inquiry statement* regarding the availability of the unit, such as "is this apartment still available?" This gives the landlord a specific question to respond to, allowing us to identify automated responses and test for differences in stated availability rates between groups. (5) is a *closing* which thanks the landlord for their time and is always followed with the applicant's full name.

Element 3, the statement of quality, is included in approximately two-thirds of all emails. The email texts which do not include a statement of quality are said to belong to the "baseline treatment." In this treatment, landlords know nothing of the applicant except their name and that they are interested in renting their apartment. When the statement of quality is included, it discloses either "positive" or "negative" information. Emails including these statements of quality belong to the positive and negative treatments respectively. Positive information always informs the landlord that the applicant has a good job and does not smoke. Negative information always informs the landlord that the applicant smokes and has a bad credit rating. These particular pieces of information were selected because they are unambiguously positive or negative. It is difficult to imagine a scenario in which a landlord would benefit from a tenant who smokes or has bad credit, two attributes which are thought to be correlated with financial irresponsibility. Likewise, it is difficult to imagine a landlord being harmed because a tenant has a good job or does not smoke.[6]

[6] Though it is possible that a landlord might infer that an applicant who makes a point of stating that they do not smoke would be turned off by an apartment building in which many tenants do smoke. It is also

This process of inquiry generation comes with a significant benefit. By pulling all texts randomly from the same pools and defining the treatments entirely by the statement of quality alone we assure that any differences in landlord responses are the result of our treatments and not the intro or inquiry texts.

Table 1B offers a summary of the number of emails sent by each applicant type as defined by their race, gender, and treatment.

B. Rental Market Data

Listings responded to were found using craigslist.org (Craigslist), an online classified ad website of enormous popularity, particularly amongst apartment seekers.[7] Craigslist is currently being viewed by 40 million unique visitors each month and is often considered one of the principal factors responsible for the fall of newspaper classified ad revenues.[8] According to a study released in March 2009 by Hitwise, an internet marketing firm, classified website visits were up 84% year over year in February 2009, and of the top 100 classifieds websites, all but three were localized sites for Craigslist (Craigslist has unique pages for each city, these are referred to as "localized sites"). According to Hitwise, approximately 2.5% (and growing) of all US Internet visits are to Craigslist. All other classified websites combined account for only 0.14% of US internet visits. Much of Craigslist's popularity is due to its offering of free posting services for landlords, employers, and those looking to sell goods and services.

---

possible that landlords could think that their apartment is "beneath" an applicant who works a good, high paying, job. To control for this, only applicants with jobs with average earnings in a certain range apply to apartments with rents in a certain percentile range for their city. This process is detailed in the appendix.

[7] When craigslist servers crashed for a matter of hours in 2009, the Christian Science Monitor ran a story under the headline, "Craigslist goes down: apartment hunting, job searching comes to a halt."

[8] According to techcrunch.com, an "internet statistics and numbers" website. A report released by AIMGroup in June of 2009 shows a fall of newspaper classified ad revenue from $16 billion in 2005 to $5 billion in 2009. During this same period, craigslist revenue has more than quintupled, from $18 million to just over $100 million.

Only one-bedroom and studio listings were sent inquiries to avoid (as much as possible) concerns of roommates, children, etc., and to make rents more comparable between any two units within an area. Many precautions were taken to avoid sending multiple inquiries in response to listings for the same apartment (these precautions along with listing characteristics and other detailed information is available in the appendix). All listing characteristics, including the address, rent, and entire text of the listing were stored in a database for analysis. Only one inquiry was sent to each listing.

Craigslist naturally lends itself to our data generation process due not only to its popularity, but also its open, text based, layout and lack of any subscription or log-in requirements. This open set up allowed us to easily record all information presented. For each inquiry we send, we are able to record: the randomly generated applicant characteristics (race, gender, treatment, and text used), landlord generated listing characteristics (location, rent, full text of ad, etc.), and landlord response characteristics (manually characterized as described below).

C. Categorization Process

The simplest response characteristic to obtain is its existence, i.e. whether or not a given inquiry receives a response. When an inquiry does receive a response, the text of the response is then manually categorized by the authors into one of several categories. To avoid experimenter bias in this categorization, all instances of applicant names (first and last, as well as email address) are automatically removed from view during categorization. The process by which this categorization was performed, along with controls taken to avoid experimenter bias and to ensure between experimenter consistency is detailed in the appendix. Broadly, responses are classified as either positive or not-positive. Positive responses state that the unit is available

and invite future contact in some manner. Non-positive responses either state that the unit is not-available, or state that the unit is available, but in a discouraging manner.[9] This categorization process is crucial to our results since we are primarily interested in the differential treatment of black and white applicants by landlords. Differences in the likelihood of simply receiving a response may be misleading since one group may receive a larger share of negative responses than the other. Our careful reading and categorization of all landlord responses takes into account the importance of the contents of a response.

IV. Results

Of the 14,018 inquiries sent, 8,980 (64.1%) received a response. Of these responses 6,556 (73%) were positive as defined in section III.[10]

A. Treatment Effects

To verify the effectiveness of our experimental design we first examine the effects of our treatments on PRR relative to the baseline (as well as the effects of our treatments on RR and PRR conditional upon receiving a response, PRR|R, are also included).[11] Table 2 reports these response rates and the significance of these pair wise comparisons, pooling all applicants regardless of race and gender.

[Table 2 – Here]

Given that the fraction of responses which were positive (PRR|R) is only 76.8% in the baseline treatment, PRR is likely to be a less noisy measure of expected applicant quality than RR. 23.2% of all responses in the baseline were negative in some way, and as such, the simple

[9] For example, a response may simply read, "Yes," or "Yes, it is available," but offer no additional information.

[10] This rather low response rate may be a result of our automated inquiries being caught up by spam filters and therefore never observed by landlords. This should be uncorrelated with race, gender, and applicant quality, and as such our results are still applicable.

[11] Again, RR = Response Rate, PRR = Positive Response Rate, and PRR|R = Positive Response Rate conditional upon receiving a response.

rate of response is likely to misrepresent true landlord beliefs about applicant quality. Put simply, considering a "no" response as equivalent to a "yes" response is likely to invite error into our data interpretation.

As shown in Table 2, applicants in our negative treatment receive a significantly lower PRR (22.1% vs. 54.0%, $p<0.001$), receiving only 41 positive responses for every 100 received by baseline applicants. Applicants in the positive treatment receive a significantly higher PRR (57.9% vs. 54.0%, $p<0.001$), receiving 107 positive responses for every 100 received by the baseline applicants. Given the significance of these differences, we conclude that our treatments were effective in their ability to manipulate landlord interest in our applicants.

That our positive treatment was ineffective in increasing overall response rates deserves explanation. This non-change in RR may have been the result of a non-effective positive treatment. However, it is difficult to argue that applicants who inform landlords that they work a well paying job are disclosing something non-positive. It could be argued that landlords simply disregarded the additional information disclosed in emails, but given the effectiveness of the negative treatment, this is not likely either. This argument is also challenged by the significant increase in PRR brought about by the positive treatment. A more compelling explanation is that landlords responded to all inquiries with an expected quality greater than or equal to that conveyed by baseline inquiries. This introduces an artificial ceiling on RR of approximately 70%.[12] As such, landlord expectation of applicant quality is still increased by the addition of positive information but this increase is only able to manifest itself in the form of an increased PRR for the positive treatment (since PRR, unlike RR, does not reach a ceiling in the baseline treatment). For these reasons we restrict our analysis to PRR.

[12] The remaining 30% would then be composed of "dead listings" (listings for units which have been rented but not removed from craigslist), as well as scams and misplaced listings not caught by our filters (as described in the appendix).

B. Black and White Name Effects

Table 3A presents the coefficients obtained from an OLS regression of race and informational treatment dummies on PRR. Column 1 contains the results for this regression, pooling males and females.

[Table 3A – Here][13]

The coefficient on Black of -0.094 is highly significant ($p<0.001$) and confirms previous findings of discrimination against African Americans. Given the intercept of 0.587, we see that, in the baseline treatment, black applicants receive about 84 positive responses for every 100 received by white applicants.

The coefficient on Positive Treatment of 0.040 is also significant ($p<0.01$), though the coefficient on the interaction term Black*Positive is indistinguishable from zero. This indicates that both black and white applicants benefit from the inclusion of positive information, though the absolute gap in PRR between black and white applicants in this treatment is unchanged. Given the small value of the coefficient on Positive, the ratio of positive responses received by blacks to positive responses received by whites in this treatment is only marginally changed (increases from 84 per 100 to 85 per 100, $p>0.10$).

The coefficient on "Negative Treatment" of -0.342 is highly significant ($p<0.001$), as is the coefficient on the interaction term Black*Negative ($p<0.01$). As a result, we observe a substantial decrease in PRR for all applicants in the presence of this treatment as well as a decrease in the absolute gap in PRR between black and white applicants in this treatment (the absolute gap falls from 9.4% in the baseline to 4.8% in the negative treatment). The ratio of PRR for black applicants to PRR for white applicants in this treatment actually *decreases* despite the

[13] A second version of this table utilizing data which removed duplicate landlords as determined by the content of their responses (i.e. their email addresses) is provided in the appendix along with a discussion concerning the costs and benefits of this additional stage of duplicate landlord prevention.

closure of the absolute gap. In the negative treatment, African Americans receive only 80 positive responses for every 100 received by white applicants.

For ease of reference, the absolute differences in PRR between black and white applicants ("gaps") as well as the ratio of black PRR to white PRR are provided in Figure 1 (along with the same figures in terms of RR).

[Figure 1 – Here]

Even though previous research has established discrimination against African Americans to be a serious and enduring problem, we were surprised to confirm these findings to the extent that we did. At the time of our data collection, U.S. apartment vacancy rates were at a 21 year high of 11.4% and the U.S. had just elected its first African American president. To find differences in landlord interest metrics of these magnitudes was unexpected, particularly since our results are likely to be understating the level of discrimination which would exist in a normal (tighter) rental market. Our results illustrate clearly the presence of the (often hidden) obstacles African Americans must face to participate in this market.

D. Gender Effects

i. Pooling by gender and examining the effect of race

Returning to Table 3A, columns 2 and 3 perform the same linear probability analysis as discussed in section B above after separating applicants by gender. The coefficients on Black for Female and Male applicants are -0.086 and -0.101 respectively (both significant with $p<0.001$). As such, we can see that in the baseline treatment, differential treatment by race is not gender specific. In the baseline, black female (male) applicants receive 85 (83) positive responses for every 100 received by their white counterparts.

The differential treatment by race for applicants of the same gender persists in both the

positive and negative treatments as well. For males in the positive treatment there is no significant change what-so-ever; neither the coefficient on Positive nor the coefficient on Positive*Black are significantly different from zero. For females in the positive treatment, however, there is a marginal increase in the relative PRR, with black females receiving 86 positive responses for every 100 received by white females in the positive treatment.

In the negative treatment, the ratio of positive responses for black males relative to white males once again experiences only a marginal change, with black male applicants receiving 84 positive responses for every 100 received by white applicants in this treatment. The effect of the negative treatment on relative response rates for females, however, is quite significant. In the negative treatment, black females receive only 66 positive responses for every 100 received by white females.

The difference in discriminatory treatment by gender seems to indicate that our negative information effects black females differently than it does black males. Given the highly randomized methodology it is extremely unlikely that female applicants were simply more likely to encounter a certain type of prejudiced landlord than were male applicants. What is more likely is that the same landlord may exhibit different types of prejudice dependent upon the characteristics of the applicant. Perhaps, for example, black females of low quality were assumed to be more likely to be single parents than were black males, white males, and white females of the same quality. If this were the case, negative information would be thought to signal something entirely different for this group than for the other three race/gender groups, and this could result in a shrinkage in the PRR gap for males, but not for females. Further investigation of this possibility would be meaningful for many people, including single-mothers.

ii. Pooling by race and examining the effect of gender

Table 4 below examines the effect of gender on PRR for black applicants, white applicants, and all applicants.

[Table 4 – Here]

Here we see that both black and white female applicants receive a larger PRR than males in all treatments, though this difference is only significant when pooling together both races of applicant and across all treatments. Pooled by race, we see a borderline significant difference of PRR in the positive treatment ($p=0.058$). Pooling across all treatments we also observe a borderline significant difference in PRR for white females relative to white males ($p=0.062$). These results are consistent with those of Ahmed and Hammarstedt (2008) who found their white female to fair significantly better in terms of landlord treatment in their correspondence study of apartment rental discrimination in Sweden.

VI. Conclusion

Despite our study taking place in a rental market facing record vacancy rates and less than a year after the U.S. elected its first African American president, we find substantial anti-black discrimination by landlords. Black apartment seekers receive, on average, 12 fewer responses per 100 received by whites, and 16 fewer positive responses per 100 received by whites. That our methodology attempted to avoid sending multiple inquiries to the same landlord makes these findings even more disturbing since they reflect the conditions of the U.S. rental market as a whole, and not the preferences of a select group of landlords. The discrimination uncovered by this study is difficult, if not impossible, to observe in one's everyday life, though it clearly affects the amount of effort African American apartment hunters must expend in order to receive the same amount of landlord attention as white apartment hunters.

Future research may wish to increase the level of positive information to identify a

change in differential treatment on the positive end. Such future research could (perhaps incrementally) increase positive information by offering references, proof of salary, specific high credit ratings, and solid rental histories.

Future research will attempt to identify the effects of neighborhood demographics on discrimination in this market. Perhaps landlords in minority neighborhoods are more accustomed to working with African American applicants and are therefore able to interpret signals sent by black applicants equal (or even lesser noise) than they do signals sent by white applicants. This would result in a decrease in the level of discrimination in all three treatments.

Another interesting avenue of research would involve analyzing our results in terms of market tightness. Landlords may find it more costly to discriminate against willing applicants when the supply of such applicants is limited. Perhaps landlords in neighborhoods with high vacancy rates are willing to set their prejudice aside for the sake of profit. a decrease in the incidence of discrimination. The behavior of large, corporate, landlords could also be examined in the context of this study.

Other interesting avenues of research include incorporating landlord specific effects, and using this methodology to examine discrimination against other minority groups, particularly the LGBT community.

I would like to acknowledge Michael Ewens and Choon Wang, who co-authored with me the first chapter of this dissertation. Their assistance was invaluable in the authoring of this chapter.

CHAPTER 1 - BIBLIOGRAPHY

Ahmed, Hammarstedt, "Discrimination in the Rental Housing Market: A field experiment on the internet," *Journal of Urban Economics* 64, pgs. 362-372, 2008

Aigner, Cain, "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review*, pgs. 175-187, 1977

Becker, "The Economics of Discrimination," *University of Chicago Press*, 1957

Bertrand, Mullainathan, "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economics Review 94 (4)*, pgs. 991-1013, 2004

Carpusor, Loges, "Rental Discrimination and Ethnicity in Names," *Journal of Applied Social Psychology 36 (4)*, pgs 934-952, 2006

Clark, "Residential segregation in American cities: A review and interpretation," *Population Research and Policy Review* 5, 1986, pgs. 95-127

Farley, Schuman, Bianchi, D. Colasanto, and S. Hatchett, "Chocolate City, vanilla suburbs: Will the trend towards racially separate communities continue?" *Social Science Research*, 7, 1978, pgs. 319-344

Fryer, Levitt, "The Causes and Consequences of Distinctly Black Names," *The Quarterly Journal of Economics 119 (3)*, pgs. 767-805, 2004

Heckman, "Detecting Discrimination," *Journal of Economics Perspectives 12 (2)*, pgs. 101-116, 1998

Heckman, Siegelman, "The Urban Institute Audit Studies: Their methods and findings," *Urban Institute Press,* 1993

Holtzer, "The Spatial Mismatch Hypothesis: What Has the Evidence Shown?" *Urban Studies*, 28, 1991, pgs. 105-122

Jencks, Mayer, "The Social Consequences of Growing Up in a Bad Neighborhood," in *Inner-City Poverty in the United States*, National Academy Press, 1990

Kern, "Racial Prejudice and Residential Segregation: The Yinger Model Revisited," *Journal of Urban Economics*, pgs. 164-172, 1981

Newburger, "Discrimination by a Profit Maximizing Real Estate Broker in Response to White Prejudice," *Journal of Urban Economics 26*, pgs. 1-19, 1989

Ondrich, Ross, Yinger, "Now you see it now you don't: Why do real estate agents withhold available houses from black customer?" *Review of Economics and Statistics 85 (4)*, pgs. 854-873, 2003

Ondrich, Stricker, Yinger, "Do real estate brokers choose to discriminate? Evidence from the 1989 Housing Discrimination Study," *Southern Economics Journal 64 (4),* pgs. 880-901, 1998

Ondrich, Stricker, Yinger, "Do Landlords Discriminate? The incidence and causes of racial discrimination in rental housing markets," *Journal of Housing Economics* 8 (3), pgs. 185-204, 1999

Page, "Racial and Ethnic Discrimination in Urban Housing Markets: Evidence from a recent audit study," *Journal of Urban Economics 38 (2)*, pgs. 183-206, 1995

Pager, "The Use of Field Experiments for Studies of Employment Discrimination: Contributions, Critiques, and Directions for the Future," *Annals of the American Academy of Political and Social Science* 609, pgs. 104-133, 2007

Phelps, "The Statistical Theory of Racism and Sexism," *American Economic Review* 62 (4), pgs. 659-661, 1972

Roychurdry, Goodman, "Evidence of Racial Discrimination in Different Dimensions of Owner-Occupied Housing Search," *Real Estate Economics 24 (2)*, pgs. 161-178, 1996

Turner, Edwards, Mikelson, "Housing Discrimination Study: Analyzing Racial and Ethnic Steering," Washington D.C. U.S. Department of Housing and Urban Development, 1991

Turner, Ross, Galster, Yinger, "Discrimination in Metropolitan Housing Markets: National Results from phase 1 HDS 2000," *Urban Institute*, 2002

Yinger, "Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act," *American Economic Review* 76 (5), pgs. 881-893, 1986

Yinger, "Acts of Discrimination: Evidence from the 1989 Housing Discrimination Study," *Journal of Housing Economics,* pgs. 318-346, 1991

Zhao, Ondrich, Yinger, "Why Do Real Estate Brokers Continue to Discriminate? Evidence from the 2000 Housing Discrimination Study," *Journal of Urban Economics 59*, pgs. 394-419, 2006

CHAPTER 1 - TABLES

Table 1A: Cities Surveyed

City Name	# of Observations	% Black	Mean Rent ($)*
Atlanta	300	61.4	758.64
Austin	196	10	718.73
Baltimore	492	64.3	847.75
Boston	1321	23.5	1062.14
Charlotte	236	32.7	725.13
Chicago	580	36.8	1089.66
Cleveland	361	51	560.92
Dallas	146	25.9	1353.91
DC	1165	60	869.58
Denver	731	11.1	729.02
Detroit	449	81.6	594.42
Houston	303	25.3	794.85
Indy	154	25.5	542.79
Jacksonville	124	29	673.80
KC	266	31.2	590.18
LA	1009	11.2	1186.72
Louisville	237	33	549.30
Memphis	110	61.4	668.02
Milwaukee	219	37.3	621.09
Minneapolis	736	18	761.13
Nashville	176	26.8	794.80
OKC	178	15.4	492.70
Philly	548	43.2	915.81
Phoenix	273	5.1	607.55
Portland	303	2	770.41
Raleigh	248	27.8	645.78
San Antonio	86	6.8	1045.72
San Jose	248	3.5	1472.56
Santa Barbara	161	1.8	612.92
SD	781	7.9	1172.81
Seattle	445	8.4	1338.36
SF	423	7.8	935.59
Tampa	664	28	678.08
Tucson	345	4.3	532.23
Average	**412**	**27**	**815.09**

*Mean Rent is Calculated Using the Rents of Units we Surveyed, not Census Data

Table 1B: Count of Observations by Race, Gender and Treatment.

Demographic	Treatment	Male	Female
Black	Negative	993	941
	Baseline	1048	1018
	Positive	1517	1479
White	Negative	969	992
	Baseline	1007	1087
	Positive	1422	1545

Table 2[++]: Overall Treatment Effects: Positive Response Rate (PRR), Response Rate (RR) and Positive Response Rate Conditional on Receiving a Response (PRR|R) Across All Treatments.

Treatment	PRR	RR	PRR|R
Baseline	0.540	0.704	0.768
	(0.008)	(0.007)	(0.008)
Positive	0.579**	0.698	0.829**
	(0.006)	(0.006)	(0.006)
Negative	0.221**	0.485**	0.456**
	(0.007)	(0.008)	(0.011)

** Statistically distinguishable from baseline treatment with $p < 0.001$ using a two-tailed t-test.
++ Positive (negative) treatment adds positive (negative) information about the applicant to baseline email text.

Table 3A: OLS Output for Demographics on Positive Response Rate (PRR)

Variable	(1) PRR (Genders Pooled)	(2) PRR (Females Only)	(3) PRR (Males Only)
Intercept	0.587***	0.591***	0.583***
	(0.011)	(0.015)	(0.016)
Black Applicant (1=Yes)	-0.094***	-0.086***	-0.101***
	(0.015)	(0.022)	(0.022)
Positive Information Treatment (1=Yes)	0.040**	0.046**	0.032
	(0.014)	(0.019)	(0.020)
Negative Information Treatment (1=Yes)	-0.342***	-0.334***	-0.350***
	(0.020)	(0.020)	(0.021)
Black*Positive	-0.001	-0.004	0.002
	(0.020)	(0.028)	(0.028)
Black*Negative	0.046**	0.027	0.064**
	(0.020)	(0.029)	(0.029)
N	14018	7062	6956
R^2	0.1021	0.1035	0.1006

** $\alpha<0.05$ ***$\alpha<0.01$

Table 3B: OLS Output for Demographics on Response Rate

Variable	RR (Females Only)	RR (Males Only)
Intercept	0.737***	0.750***
	(0.014)	(0.014)
Black Applicant (1=Yes)	-0.059***	-0.099***
	(0.020)	(0.020)
Positive Information Treatment (1=Yes)	0.005	-0.013
	(0.017)	(0.018)
Negative Information Treatment (1=Yes)	-0.192***	-0.244***
	(0.021)	(0.021)
Black*Positive	-0.017	0.011
	(0.026)	(0.026)
Black*Negative	-0.055*	0.053*
	(0.030)	(0.030)
N	7062	6956
R^2	0.0488	0.0467

* α<0.10 ** α<0.05 ***α<0.01

Table 4: The Effect of Gender on Positive Response Rates (PRR) by Race and Treatment

	Baseline	Positive	Negative	Pooled
All	F: 54.9	F: 43.9	F: 22.9	F: 48.0
	M: 53.1	M: 41.7	M: 21.5	M: 45.6
	p = 0.250	p = 0.058	p = 0.290	p = 0.004
White Only	F: 59.1	F: 63.7	F: 25.7	F: 51.9
	M: 58.3	M: 61.5	M: 23.3	M: 49.7
	p = 0.721	p = 0.225	p = 0.220	p = 0.062
Black Only	F: 50.5	F: 54.7	F: 19.9	F: 43.9
	M: 48.2	M: 51.6	M: 19.6	M: 41.7
	p = 0.295	p = 0.090	p = 0.897	p = 0.058

F – PRR for Females M – PRR for Males p – p-value for t-test of difference in F and M

CHAPTER 1 - FIGURES

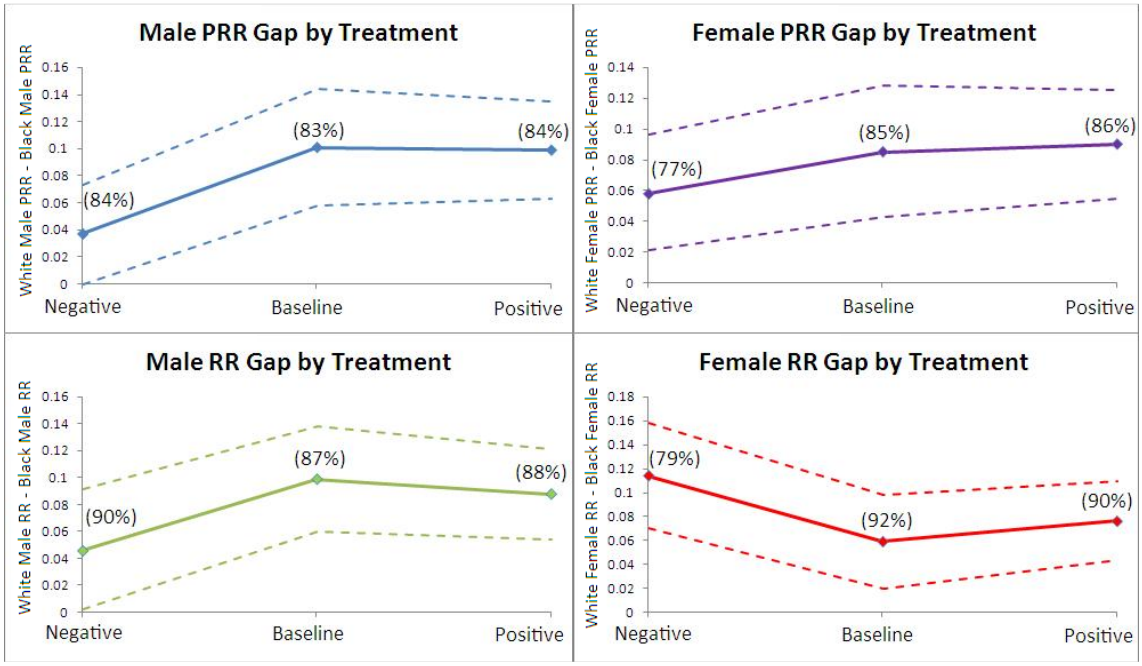

Figure 1: Absolute Differences (Gaps) in Black and White Response Rates (RR) and Positive Response Rates (PRR) for each treatment (Negative, Baseline, and Positive), by gender.* Ratios of Black/White RR and PRR Given in Parenthesis**

***95% confidence intervals of these gaps are dotted.**

****** For example, in the negative treatment for males, white male applicants are 4% more likely to receive a positive response, and black males receive 84 positive responses for every 100 received by white males.**

CHAPTER 1 – APPENDIX A – OTHER RESULTS OF INTEREST

Neither the response rate, nor the gap in response rates between white and black applicants differs significantly for units which are more expensive relative to other units in the same city. As shown in Figure 2, both measures are rather consistent across rental prices. These steady response rates may reflect the fact that landlords with higher priced units have not only the most to gain by renting a unit quickly, but also the most to lose by renting to a bad tenant. That the RR gap between black and white applicants is statistically no different for units with rents below the 10^{th} percentile than it is for units with rents above the 90^{th} percentile is telling. Supposing that the rental price of a unit reflects the typical rental price of an equivalent unit in the same neighborhood, landlords do not appear to be steering black applicants away from expensive neighborhoods or towards cheaper neighborhoods. This deduction, however, is likely to be flawed. A more thorough dissection of rent effects would take into account neighborhood specific rental information, in particular, local vacancy rates.

Another interesting finding of our study is that the percentage of residents who are black within the city in which an apartment is located has no significant effect upon the incidence of discrimination. As shown in Figure 4, the percentage of blacks living, if anything, only seems to increase the degree of differential treatment, though this effect is not statistically differentiable from zero. Figure 5 offers city specific differences in RR and PRR, as well as the number of observations we have for each city, and the percentage of black residents within each city. A city level breakdown of such effects is obviously flawed since cities with a large percentage of minority residents may actually be very segregated (i.e. the demographics of Chicago's South side are quite different than those of Chicago's North side). A better examination of this effect would use neighborhood level demographic information instead.

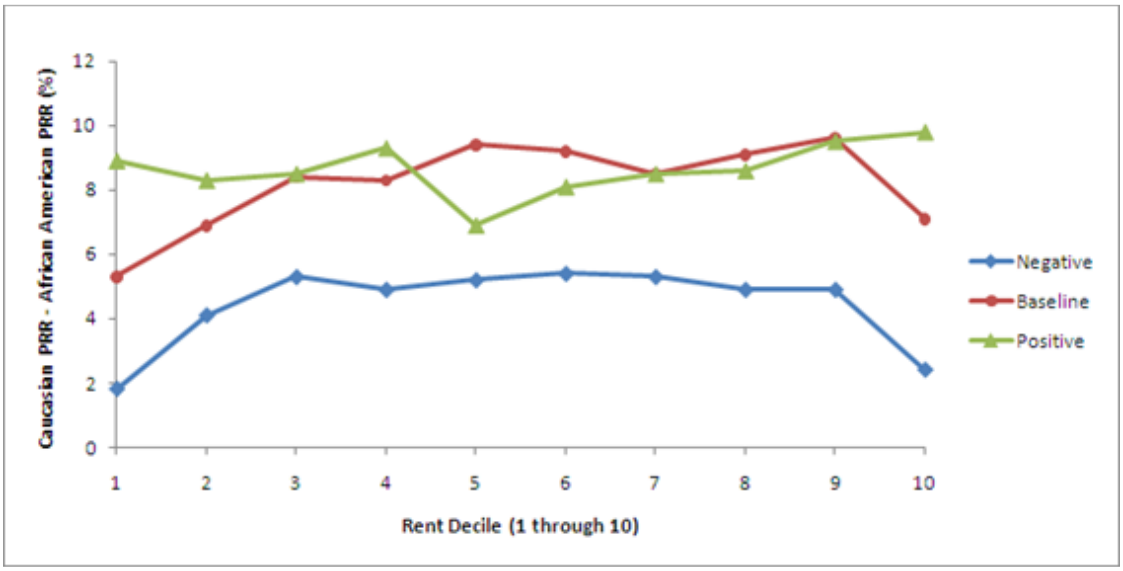

Figure 2: Positive Response Rates (PRR) Gaps by Apartment Rent

Rent Decile is determined by city. So, data in the 1st decile is composed of all apartments with rents amongst the lowest ten percent for the city in which they are located. Each data-point therefore contains observations from 10% of the units surveyed in each city for each treatment.

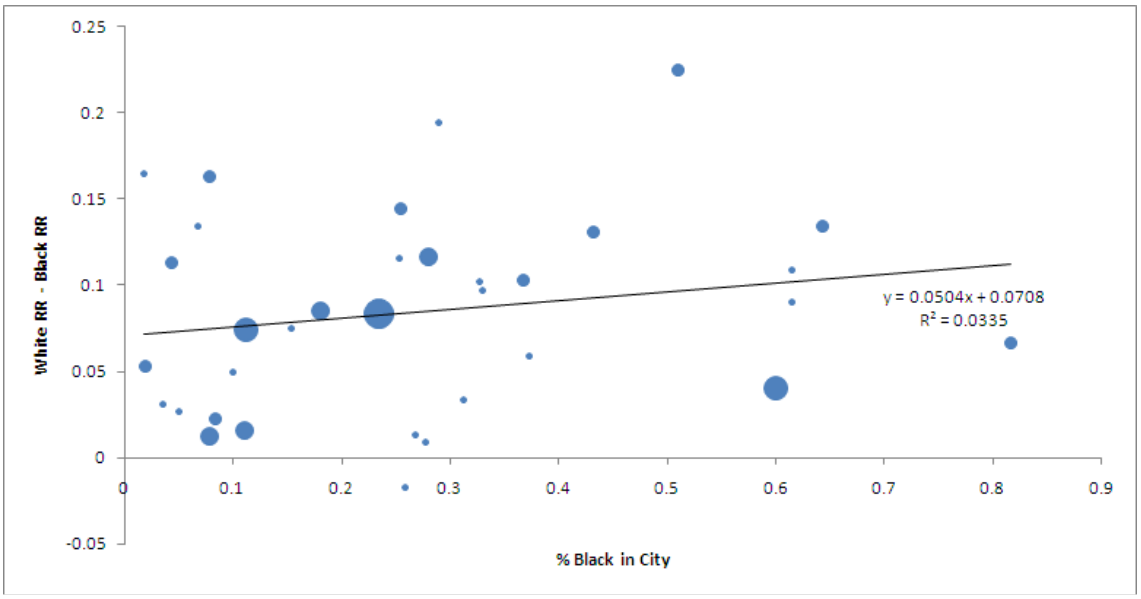

Figure 3A: Differences in Black and White Response Rates (RR) by % Black in City: Dot Size Reflects Number of Observations for a Given City

Slope coefficient not statistically distinguishable from zero (p=0.300).

Regression of %Black on Black – White Positive Response Rate (PRR) also returns a slope coefficient statistically indistinguishable from zero (0.0379, p=0.408)

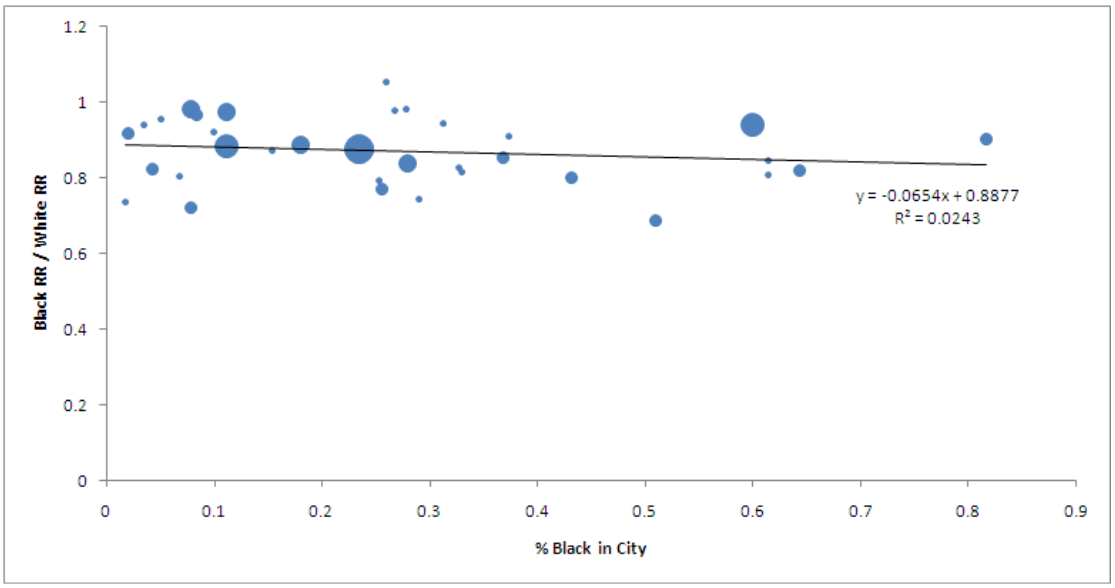

Figure 3B: Ratio of Black/White Response Rate by % Black in City: Dot Size Reflects Number of Observations in Given City

Slope coefficient not statistically distinguishable from zero (p=0.388).
Regression of %Black on Black – White Positive Response Rate (PRR) also returns a slope coefficient statistically indistinguishable from zero (-0.064, p=0.497)

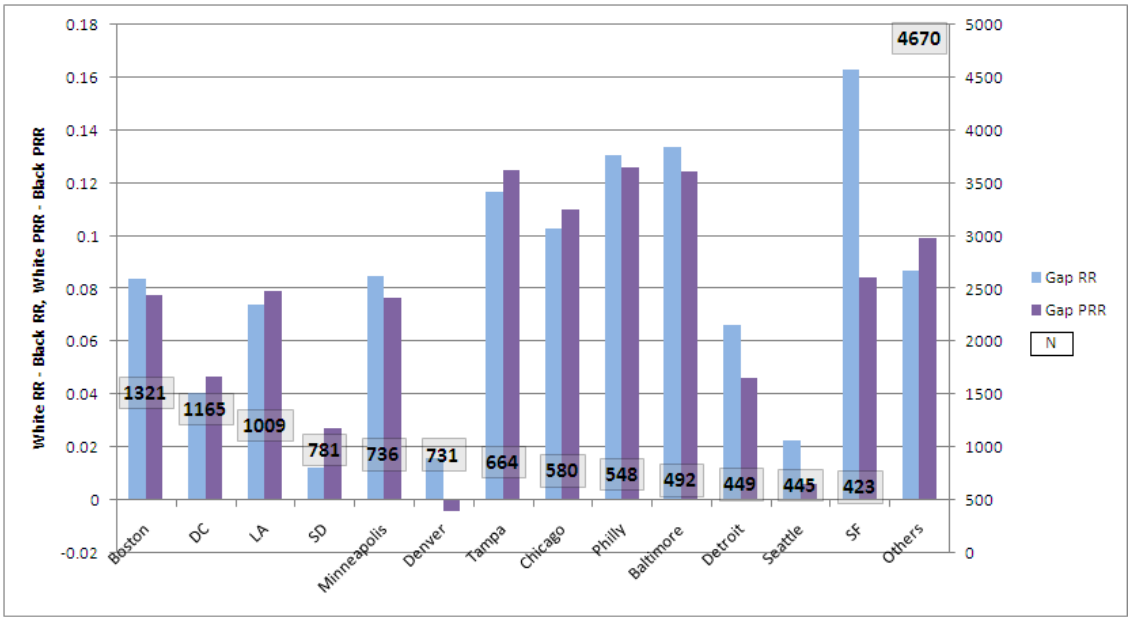

Figure 4A: Differences in Black and White Response (Positive) Response Rates (PRR) Observations, N, by city included. Individual cities with N > 400, Others grouped

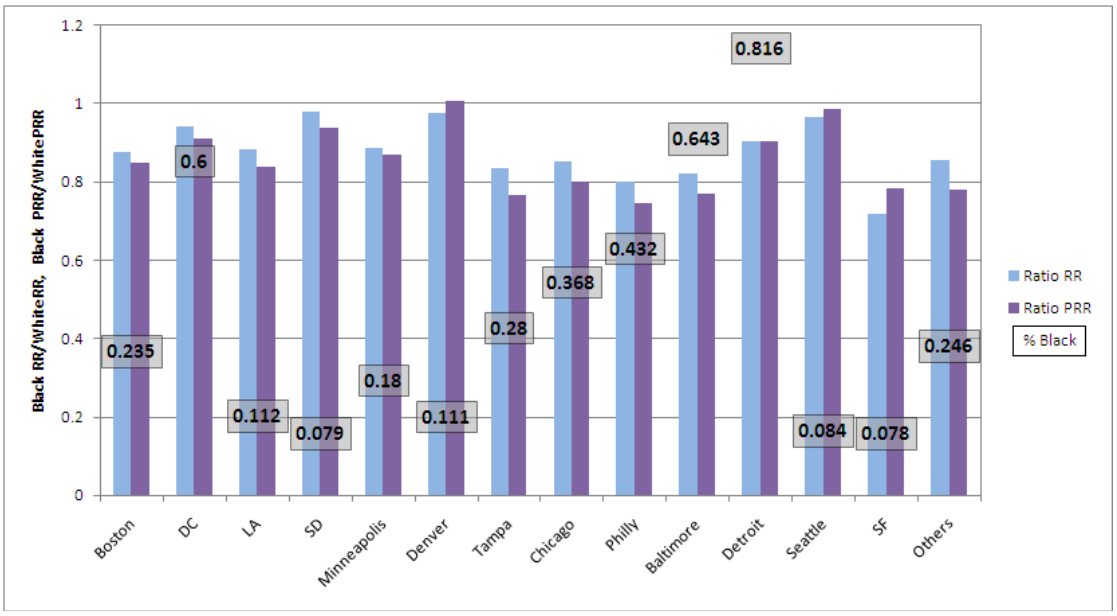

Figure 4B: Ratio of Black/White Response Rates (Positive) Response Rates by city. % Black by city also included (Boxed). Individual cities with N > 400, Others grouped

CHAPTER 1 – APPENDIX B – RESULTS, DROPPING DUPLICATE LANDLORDS

As explained in Section III and detailed in Appendix C (detailed methodology), several measures were taken to avoid sending multiple inquiries to the same landlord. This was done for two reasons: First, many landlords post multiple listings for the same apartment and we did not want to harass them with multiple inquiries, or expose our study to them by sending them two (or more) identical emails with different names (this event is highly unlikely given our experimental design). Second, we are attempting to identify small landlord behavior and avoid sending to large apartment complexes. The data presented in the body of the paper utilizes data gathered using these processes.

The data presented in this portion of the appendix utilizes data obtained after a second stage of duplicate landlord removal which screens landlord responses rather than listings. We check all responses for unique strings defining landlord emails and/or phone numbers then eliminate every response containing a duplicate string after the first. We also disregard the inquiries to which these responses were sent as though they had never been sent. As such, if a landlord (as defined by a unique email address or phone number) responds to more than one inquiry, all responses after the first are removed from our sample, as are the inquiries which generated these responses.

The data generated by this process is superior in regards to its increased assurance that we are only using one response from each landlord. The data is inferior due to the fact that it introduces bias into our sample in the following way. Suppose a landlord only responds to high quality applicants (never responding to low quality applicants). Were this landlord to receive three emails from baseline or low-quality applicants, we will observe three inquiries sent with no responses received. Were this same landlord to instead receive three high quality emails, we

would observe three inquiries and three responses, but would then drop the last two. When the landlord responds, they are treated as one landlord, but when they do not, they are treated as three landlords. The result is an over-emphasis on landlords who only respond based upon some criteria (i.e. only respond to white applicants, or high-quality applicants).

Our original screening process of duplicate landlords – which does not introduce any bias – is quite robust. Dropping duplicate landlords in this second stage only reduces our sample by 427 observations (3%). The results of the OLS regressions presented in Table 3 are re-produced below with this new sample. (Table on next page)

The only significant change resulting from this altered sample is that the black*negative interaction coefficient on RR for women falls out of significance (p=0.111, as opposed to p=0.068). Why this happens is not entirely clear and will not be speculated upon at this time.

Table B1: OLS Output for Demographics on (Positive) Response Rate

Variable	RR (female)	PRR (female)	RR (male)	PRR (male)
Intercept	0.729***	0.583***	0.740***	0.571***
	(0.014)	(0.015)	(0.014)	(0.016)
Black Applicant[+] (1=Yes)	-0.064***	-0.086***	-0.098***	-0.096***
	(0.020)	(0.022)	(0.021)	(0.022)
Positive Treatment (1=Yes)	0.002	0.043**	-0.014	0.040*
	(0.018)	(0.019)	(0.018)	(0.021)
Negative Treatment (1=Yes)	-0.198***	-0.333***	-0.247***	-0.339***
	(0.021)	(0.021)	(0.022)	(0.021)
Black*Positive	-0.014	0.000	0.010	0.006
	(0.027)	(0.028)	(0.027)	(0.028)
Black*Negative	-0.049	0.029	0.052*	0.054*
	(0.031)	(0.029)	(0.031)	(0.029)
N	6834	6834	6757	6757
R^2	0.0492	0.1032	0.0474	0.0990

* $\alpha<0.10$ ** $\alpha<0.05$ ***$\alpha<0.01$

CHAPTER 1 – APPENDIX C – DETAILED METHODOLOGY

Of course, any massive automated process is bound to be accompanied with natural shortcomings. Multiple safeguards were put in place to avoid any foreseeable issues. Before introducing our detailed methodology we will first discuss the three most prevalent shortcomings we encountered:

Shortcomings of Methodology

One: Craigslist offers landlords the opportunity to mask their contact email with an anonymous forwarding address of the form "*[random-number]@craigslist.org*." As a result of this, we are unable to ensure that we do not send multiple inquiries to the same landlord. We take precautions against this (see detailed methodology below), but cannot eliminate this with certainty.

Two: given the volume of emails being sent out from our servers, some of our inquiries may have been automatically redirected to the recipient's spam folder. This, however, should be uncorrelated with the race and gender of the applicant.

Three: while many controls were taken to ensure consistency there still existed room for discrepancy between categorizers. Since experimenters were unable to see the name of the applicant receiving the response, such discrepancies should be randomly distributed amongst all responses and therefore uncorrelated with race and gender. However, if a certain type of response was only sent to applicants of a certain race and/or gender, and these types of responses were categorized differently by the experimenters, this would lead to a misrepresentation of the actual distribution of responses.

Detailed Methodology

Part C of this section contains a succinct overview of the methodology for those readers simply interested in the overall data generation/gathering process. The methodology and reasoning behind every step of this process is detailed in sections A, B, and D.

A. Applicant Generation

An "applicant," is defined as the combination of a randomly determined: gender, ethnicity, email text. This section discusses the creation of each.

Name Generation

This being a correspondence study, the only means of denoting an applicant's race to a landlord is through their name. It is therefore imperative that names be carefully selected to create a "perceived ethnicity" in the mind of the landlord.[14] Potential first and last names were gathered by pooling:

1. All names (first and last) utilized by BM2005.
2. Last names of high frequency amongst blacks and whites, provided by the US Census name survey.

Table C1a: First Names Used in Experiment

White Female	White Male	Black Female	Black Male
Allison	Brad	Aisha	Darnell
Anne	Brendan	Ebony	Hakim
Carrie	Geoffrey	Keisha	Jamal
Emily	Greg	Kenya	Jermaine
Jill	Brett	Lakisha	Kareem
Laurie	Jay	Latonya	Leroy
Kristen	Matthew	Latoya	Rasheed
Meredith	Neil	Tamika	Tremayne
Sarah	Todd	Tanisha	Tyrone

[14] In the language of BM2005, an applicant must have a sufficiently "white sounding" or "black sounding" name.

Table C1b: Last Names Used in Experiment and Percentage of U.S. Population, by race, with Each Last Name

White Last Name	% White with Name	% Black with Name	African American Last Name	% White with Name	% Black with Name
Bauer	96.34	0.34	Washington	5.16	89.87
Becker	96.4	0.46	Jefferson	18.72	75.24
Erickson	96.39	0.24	Booker	30.09	65.57
Klein	96.53	0.3	Banks	41.3	54.24
Kramer	96.47	0.28	Mosley	42.68	52.83
Mueller	96.96	0.23			
Schneider	96.67	0.33			
Schroeder	96.74	0.23			
Schwartz	96.77	0.4			

The African American last names utilized by BM2005 were: Jackson, Jones, Robinson, Washington, and Williams, only one of which (Washington, also used in our study) is held by a larger percentage of African Americans than white Americans.

Email Address Generation

Each applicant, once assigned a name, was assigned a unique email address of the form: Firstname.lastname@domain.com where the domain was privately purchased for the sake of this research and served as the host for the email addresses. Eleven domains were purchased. The names of these domains were selected so as to not indicate any sort of quality. The domain names used were: gophercontact.com, histelect.com, inertcom.com, kwik-email.com, mailtymes.com, reelofmail.com, tigersure.com, timefourmail.com, urmailplace.com, votolom.com, wiredispatch.com, yourmailspace.com

As a precaution, we created generic “this page under construction” websites for the domains and actively tracked traffic to these sites. Significant traffic would have suggested that landlords receiving our emails found the domain names suspect. Analytics on traffic combined with the mapping of domain to city allows us to gauge whether a certain set of landlords are viewing our domains after receiving our emails. These analytics show a minute amount of traffic to our sites which does not correlate with the cities in which these landlords should be located.

We therefore conclude that our domains are neither suspicious nor in any way harmful to our study.

Email Content Generation

The content of each email was broken into the following six "parts":

Table C2: Email Composition

Part of Email	Description
Subject	The subject line of the email
Salutation	The basic one word opening.
Introduction / statement of interest	An introductory statement containing the applicant's name, a statement of interest, and a clarification of the particular apartment being inquired about.[15]
Inquiry	A question concerning whether the apartment was still available.
Body	Information about the applicant quality: positive, negative, or omitted.
Closing	A simple closing such as "thank you."
Signature	The full name of the applicant.

Emails took the following generalized format:

[Salutation],
[Introduction and statement of interest]. [Body (if necessary)]. [Inquiry].
[Closing],
[Signature]

Given this format the applicant's name appeared three times, in their email address, in the introduction, and in the signature. It is therefore highly unlikely that the landlord did not perceive the name of the applicant. An example inquiry is offered below:

> Hello,
>
> My name is [Full Name], and I am writing in response to your listing for an apartment for [apartment rent]/month. I am a software engineer and I do not smoke.[16] Is this unit still available? Thank you for your time,
>
> [Full Name]

[15] This clarification took the form of a statement of the apartment's rent. This also served to make the emails appear more authentic.

[16] Had this been assigned negative information, this sentence may have instead read: Just so you know, I smoke and my credit rating is not very good.
Had this been assigned to the control group, this sentence would have been omitted entirely.

Body Generation

The body of each email contained information concerning the quality of the applicant. Since this research is aimed at identifying statistical discrimination, it was necessary that the amount of information contained in the body was varied in a controlled manner, and in a way perceivable to the landlord.

Table C3: Email body information

Quality of Information	# of pieces of information	Specific pieces of information presented in the email body
None	0	None
Positive	2	Non-smoker. Good job.
Negative	2	Smoker. Bad credit.

The information pieces presented are all easily stated and intuitively seen to be positive or negative. Informal interviews with landlords and apartment seekers concluded that these pieces of information were commonly stated in early communications with a landlord. The occupations for "good job" were sorted into three earnings categories using 2006 data obtained by the Bureau of Labor Statistics[17]. Applicants were randomly assigned an occupation which was relatively "Low-paying," "High-paying" or "Neither low-nor-high paying." Those with low paying jobs (relative to the other occupations) were only eligible to apply for apartments in the 20th to 80th percentile of rents. Those with high paying jobs were only eligible to apply for apartments in the 40th to 90th percentile of rents. Those with neither low-nor-high paying jobs were eligible to apply for any apartment within the 20th to 90th percentile of rents. The occupations used were:

[17] www.bls.gov

Table C4: Occupations

Occupation	Broad Category
accountant	financial specialists
civil engineer	architecture and engineering occupations
computer programmer	computer and mathematical occupations
computer software engineer	computer and mathematical occupations
dental assistant	Healthcare Support Occupations
designer	arts, design, entertainment, sports, and media occupations
environmental engineer	architecture and engineering occupations
financial analyst	financial specialists
human resource manager	Management occupations
insurance sales agent	Sales and Office Occupations
legal assistant	legal occupations
market researcher	life, physical, and social science occupations
medical assistant	Healthcare Support Occupations
pharmacist	healthcare practitioners and technical occupations
physical therapist	healthcare practitioners and technical occupations
purchasing agent	Business operations specialists
sales manager	Management occupations
sales representative	Sales and Office Occupations
salesperson	Sales and Office Occupations
system analyst	computer and mathematical occupations

B. Apartment Selection

All apartment listing replied to were found through craigslist.org, an online classified ad site organized by city. The primary goal of our apartment selection algorithm was to weed out scams, misplaced listings, repeated listings, and listings posted by individuals with "non-landlord" incentives. Those with non-landlord incentives include private "apartment finders" who make a living as middlemen between landlords and renters, existing tenants looking for someone to take over their lease, and landlords requesting a certain gender of tenant (this gender seeking behavior is actually prohibited by craigslist.org, though it still occurs regularly enough to require a check).

This section details the criteria used to select the listings to which the applicants would write inquiries. Ultimately, only 3.56% of all of all craigslist apartment listings in the surveyed cities were sent an inquiry.

Number of Bedrooms

Only one-bedroom and studio apartments were considered. This increased the utility of rent as an indicator of quality substantially (four bedroom apartments renting for $1000/month are of substantially lower quality than one bedroom apartments at the same rent). This also avoided potential issues associated with landlord assumptions. In particular, we wanted to avoid sending a signal which indicated that the applicant was more likely to be a single mother (i.e. a female applying for a 2+ bedroom apartment) or have roommates. In general, this helped eschew landlord assumptions about the number of people who would be residing in the apartment and why these other tenants were not sending the inquiries. 56% of the apartment ads posted on craigslist were not sent an application due to this criteria.

Rent Selection

Over the course of a month, data was collected on the rental prices of all one-bedroom and studio apartments listed on craigslist.org for each city. Within each city, units with rent below the 20^{th} or above the 90^{th} percentiles were dropped from the sample. This was done entirely to avoid misplaced postings. On the low end we eliminate postings for parking spaces, storage lockers, individual rooms, trailer hook-ups, and weekly rentals, which have separate portions of the website devoted to them but are commonly mis-categorized. Also eliminated on the low end are misleading rents, such as “$1 for the first month” or “half off the first month” which landlords often indicate as the true rent in the interest of meeting more potential tenants’ search criteria. On the high end we rid of postings for houses for sale and mortgages.

Also dropped were ads which did not include the rental price of the property (7.1% of all ads), and ads which had the word "free" in the title of the posting (2.1% of all ads), since the majority of these ads offer a certain number of month's rent "free" but then incorporate this discount into the actual rent over the terms of the lease, double counting the discount and providing a misleading rental price.

Specific Listing Criteria

Listings containing certain terms were dropped to avoid sending multiple inquiries to the same landlord, and to avoid those with "non-landlord" incentives.

Table C5: Dropping Criteria

Dropped if contains…	**Reason for Dropping**
"Finder*", "Hunter*", "Seeker*", "Locater*"	These indicate the poster is a professional apartment finder, not a landlord.
Very large jpeg	To circumvent craigslist.org's repeat-posting rules/checks, some landlords post large images of their ad instead of text. We are unable to record data from these images.
"a week", "per week", "/week", "weekly"	Weekly rentals are commonly short term, and certainly misrepresent the rent as 7/30 or 7/31 of the true monthly rent.
"female"	Listings rarely contain the word female unless the landlord is seeking female tenants only.
"free" (in title)	Offer a certain number of month's rent "free" but then incorporate this discount into the actual rent over the terms of the lease, double counting the discount and providing a misleading rental price.
Profanity	Listings containing profanity were not sent applications. This avoided scams and pranks.

*Allows for multiple endings, such as "finders"

Duplicate Prevention Criteria

To avoid sending multiple applications to the same landlord, the following safeguards were utilized:

1. Search for phone numbers within each listing. If the same phone number is in multiple listings, only apply to one of them.

2. Search for links to external websites within each listing. If the same link is included on multiple listings, only apply to one of them.
3. Search for images within each listing. If the same image is in multiple listings, only apply to one of them.

Only listings which provided an email contact were applied to. Those without an email contact were dropped (3.7% of all ads – the vast majority of landlords provide an email contact). Listings with the same email contact were only applied to once.

Timing

To avoid "spamming" landlords, we did not send applications to every ad which met our criteria. Only ads older than 1 hour and younger than 48 hours were sent applications. The lower limit on the age prevented us from appearing automated by writing to listings the moment they were posted. The upper limit decreased the likelihood of the apartment already being taken, potentially increasing RR.

Applications were assigned to ads according to the following algorithm (based upon the time the application was created):

35% of the time, assign the application to an ad with age between 1 hour and 4 hours.

20% of the time, assign the application to an ad with age between 4 hours and 6 hours.

45% of the time, assign the application to an ad with age between 6 hours and 48 hours.

The same listing never received multiple inquiries.

C. Sending/Responding Process

Given the above setup, the sending process is broken down into the following eight steps:

1. Randomly select applicant gender
2. Randomly select applicant race

3. Randomly select a name of the appropriate gender and race
4. Assign the corresponding email address at a random domain
5. Randomly select an amount of info (none, positive, negative)
6. If info == positive, randomly assign a job
7. Generate email text
8. Randomly assign a listing and send email (recording all listing information)

Should a landlord respond to the sent email, the following three step process is then followed:

1. Record the text of the response and the time it was received.
2. Respond no sooner than 1 hour and no later than 2 hours afterwards with an automated response to avoid wasting the landlord's time.[5]
3. Hold landlord response data for categorization. Delete landlord email address and name in accordance with human subjects protection protocol.

Application emails were sent out Monday through Saturday between the hours of 9 am – 12 pm, 1 pm – 4 pm, or 6 pm – 9 pm, over the course of 8 weeks. Each server sent one email every 12 minutes during these times.

D. Landlord response categorization

Emails sent by landlords in response to the randomly generated inquiries were manually categorized by the authors. Two authors categorized each response according to the criteria listed below (Table C7). To ensure responses were consistently categorized, the authors pre-categorized 1,000 responses each, discussing their methodology and reasoning for each categorization and making sure the other would have categorized the response in the same manner. Afterwards, these 2,000 responses were de-categorized and then re-categorized to avoid any learning or ordering effects.

Table C6: Response Categories

Category	Description
Available	The apartment is unambiguously stated as being available *and* future interaction is encouraged, i.e. a showing time is proposed or requested, they ask for future emails/phone-calls, etc.
Not Available	The apartment is said to be not available (unavailable), but no reason is provided as to why.
Not Available + reason	The apartment is said to be unavailable and a reason is given. The most common reason is that the apartment has already been rented.
Ambiguous leaning Yes	It is not clearly stated whether the apartment is available, but the language seems to indicate it is. i.e. “Thank you for your email. Feel free to call me whenever you like.”
Ambiguous leaning No	It is not stated whether or not the apartment is available, but the language seems to indicate it is not. i.e. “We may have other properties you are interested in become available.”
Available but disinterested	The landlord states the apartment is available but does not attempt to promote future contact/interaction. i.e. [Start of email] “The apartment is available.” [End of email].
Available + requirements	If any of the requirements were discussed/restated, such as: income, credit score, single resident only, no pets, full deposit, lease restrictions, etc.
Available + if	The unit is technically available, but an application has been submitted and the unit will only be available if this application falls through.
Available + more info	If the landlord requested more information concerning the quality of the tenant (i.e. not simply for their phone number): income, credit, number of residents, type of job, pets, etc.
Scam	A response which is clearly an attempt to obtain money or valuable information from the applicant.
Auto-reply	An “out of the office” auto reply that cannot be interpreted as any human response.

CHAPTER 2

THE IMPORTANCE OF APPEARING EARNEST: ALTRUISTIC BEHAVIOR IN THE PRESENCE OF MISUNDERSTANDINGS

Bryan Tomlin

UCSD Department of Economics

MAY 2010

Abstract

Using a unique experimental design which creates a dichotomy between how a subject behaves and how they are perceived as behaving, we test whether altruistic behavior observed in the laboratory is the result of preferences over certain outcomes or preferences for appearing to have created certain outcomes. Taking advantage of a within subject design, we add greater nuance to our understanding of altruistic behavior and egalitarian preferences, identifying specific and diverse motives for giving within each subject. Despite idiosyncrasies in between subject behavior, subjects can be grouped into one of three general types: "Self-interested" subjects who behave selfishly regardless of how their behavior is perceived by others. "Appearance-concerned" subjects who consistently exploit informational asymmetries to obtain self-biased earnings distributions while maintaining a socially favorable appearance. And "True-egalitarians" who divide earnings equally, even if doing so makes them appear selfish. The combined behavior of these three groups replicates previous dictator game findings.

I. Introduction

Suppose you are ordering a cup of coffee from your usual barista and, for whatever reason, you are interested in giving him a tip. Further, suppose that you have a one-dollar bill and two quarters and that you would prefer tipping the barista one dollar. The problem is, you know he will be turning around to make your coffee and will likely not perceive you placing a dollar bill into the tip jar; though he will certainly be able to hear your two quarters hitting the jar. As a result, you throw in the two quarters, causing the barista to turn and say "thank you." In the interest of being perceived as more altruistic, you behaved less altruistically.[18]

This paper presents an experiment designed to test the above mentioned phenomenon, namely, the effect of recognition on altruism. A derivative of the standard dictator game is presented which offers subjects the opportunity to give money away in a scenario where they know in advance that greater generosity may be perceived as being less generous and lesser generosity may be perceived as being more generous. An emphasis is placed on within subject behavioral identification, allowing the relative effects of selfishness, efficiency, fairness, and recognition to be tested for each subject individually.

Behavior is highly heterogeneous but highly consistent within the majority of subjects. Generalizing this behavior for the sake of summary, we find that subjects can be broken down into three broad categories: those who are self-interested and are content to be observed as such, those who are strongly influenced by how they will be perceived by others (these subjects will exploit informational asymmetries to behave as selfishly as possible while preserving an image of being "fair"), and finally, those who will behave in an equitable manner regardless of

[18] You certainly could have thrown in the whole $1.50, but we'll assume your self-interest kept you from doing that.

the way they are perceived. This latter group is often, but not always, highly concerned with efficiency as well.[19]

Section II discusses the ways in which this paper builds upon the existing literature. Section III presents the general experimental procedure used to create the desirable level of "social distance" between dictators and recipients. Section IV presents the methodology of each treatment and analyzes the results. Section V concludes.

II. Previous Work

In a standard Dictator Game (DG), one player, the Dictator (D), is given some amount of money (m) and the opportunity to pass any amount of it ($x \leq m$) to her partner, the Recipient (R). Despite the lack of a strategic monetary incentive for D to pass $x > 0$, since the advent of this experiment (Forsythe et al. 1994) mean offers of 30% created by the 70% of D's choosing to share non-zero amounts have been the observed norm.

A variety of explanations have been proposed to justify this seemingly selfless behavior. Andreoni (1990) suggested that subjects may experience a "warm-glow" from giving, positing that dictator utility is in part a function of recipient earnings. Rabin (1993) and Fehr and Schmidt (1999) propose models of "fairness" preferences which suggest that subjects gain utility (receive disutility) from equal (unequal) earnings. Andreoni and Vesterlund (2001) show that efficiency motives may be responsible for altruistic dictator behavior in a modified dictator game in which social earnings are not fixed.

Most relevant to this study are examinations of giving which posit that recognition effects are responsible for the observed selfless behavior. Hoffman et al. (1996) (hereafter HMS) examined the effect of recognition on giving by altering the level of experimenter and recipient blindness across several treatments of a standard dictator game. HMS found that the number of

[19] Not in the sense of Pareto efficiency, but rather social payoff maximization.

purely selfish (x=0) offers was increasing in anonymity. Bohnet and Frey 1999 (BF) utilized several between subject dictator game treatments, one in which D and R were anonymous to one another, a second where D first viewed R but R never saw D, and a third where D and R each viewed the other (all encounters were face to face) to study this same effect. BF found a highly significant increase in the number of egalitarian (x = m/2) offers as the level of social distance decreased. Specifically, when D and R each viewed the other 71% of D chose x = m/2 (SE = 2.43%), relative to only 25% when D and R were anonymous to one another.[20] Neither HMS nor BF can determine whether the increasing number of egalitarian distributions associated with decreasing social distance is a result of a complimentary nature between other-regarding-allocations and the recognition one receives for creating such allocations, or if some D's simply cannot stand to be viewed as selfish and must behave selflessly to prevent this image. This paper builds upon the work of HMS and BF and tries to answer this question by creating a disconnect between true behavior and perceived behavior.

Other studies have identified the effect of recognition on other regarding behavior by utilizing endogenous choice methods. One such study is presented by Andreoni and Bernheim 2007 (AB). Here, the allocation chosen by D is randomly over-ruled by nature with probability p (known to D), and replaced with x=0 or x=m. AB find that, as p increases, the percentage of D's offering x=0 increases: as the probability that nature was responsible for the selfish outcome increases, the number of D's choosing to behave selfishly also increases. In another study, Lazear, Melmendier, and Weber 2008 (LMW) allow D to quietly opt-out of the DG, keeping some amount while keeping R ignorant of the fact that a DG was ever on the table. By increasing the size of m within dictators – holding the amount D will keep should they opt out constant – LMW

[20] The level of interaction in the experiment presented by this paper most closely resembles the face to face interaction treatment of BF; indeed, the baseline DG performed by this study presents data closely resembling that found by BF in their face-to-face treatment.

is able to identify three different types of dictator: those who enjoy sharing, those who dis-enjoy not-sharing, and those who do not like to share. The research presented by this paper also identifies these three groups, but builds upon LMW's and AB's findings by allowing D's to actually appear less-other-regarding while in reality being more generous. This allows us understand in greater detail those who LMW would classify as those who enjoy sharing. In LMW, this group always opts into the DG and shares a significant amount – it is therefore impossible for them to determine whether these individuals prefer *being* generous or *being perceived* as generous (since they cannot have one without the other). This study creates the disconnect between true and perceived behavior necessary to determine that this group, in fact, enjoys being generous – regardless of how their behavior is perceived.

III. Experimental Procedure

A. Ensuring Experimenter Blindness and Dictator Culpability

The experiment utilized by this study is a variation of the canonical dictator game (DG). A baseline DG was also run within subjects to test for consistency with previous experiments. The entire experiment required approximately 90 minutes. 62 subjects were used[21], all of whom were undergraduates at the University of California, San Diego, recruited using fliers posted at various locations throughout the campus.

Divisions were made in the form of "tokens" each worth ten cents ($0.10). The experiment consisted of 10 "settings" (allocations), as well as a baseline DG. At the end of all 11

[21] 6 of these subjects were "monitors" who did not participate in the experiment directly, but instead ensured that every stage of the process proceeded as the experimenter said it would. Monitors were considered necessary due to the splitting up of subjects into two groups in two separate rooms – monitors ensured subjects that both groups were doing what it was said they would be doing and that subjects were not simply dismissed. Monitors were paid a flat rate of $20. This amount was given to them privately and the amount they received was not disclosed to other subjects to avoid any sort of bias this may have caused (i.e. feeling entitled to $20 themselves).

settings, one setting was selected at random to be carried out for actual payments.[22] Subjects earned an average of $21.80.

Since our goal was to study the effect of perception on altruism it was necessary to utilize an experimenter blind format. It was also necessary that the dictators feel "responsible" in the eyes of the recipient. Finally, given that informational asymmetries exist between the dictator and the recipient, it was also necessary that the dictators did not believe that they would be able to contact their paired recipient after the experiment in order to explain the actual procedure of the experiment. To accomplish this, the following experimental procedure was used:

1. Subjects were run in groups of 10 (or 8 depending on attrition), and were asked if they knew anyone else in the room. If they did, both parties were dismissed.
2. Subjects were randomly separated into two groups of 5 (or 4) – Group A (A), and Group B (B), dictators and recipients respectively.
3. Subjects were informed that they would each be paired with a single partner from the other group, and that the B's would know who their partner was, but the A's would not. They were also informed that the A's would be solely responsible for determining the payments for both themselves and their partner.
4. A's were given an ID#, 1 through 5, and asked to stand and introduce themselves by their ID#. B's were told to pay attention since they would later be assigned to a partner using these ID #s.

[22] Subjects also completed nine additional allocations of another altered dictator game. The results of these nine settings are not presented here. This second experimental condition was run *after* the experiment described in this paper and, as such, it will be assumed that the second experimental condition had no effect upon the results presented by this paper.

5. B's were lead into a separate room by a separate experimenter and randomly assigned a partner (using the ID #s). B's would never see their partner or the original experimenter again. A's would never see the second experimenter, and would never see their partner again. A's were informed that this second experimenter, who would only know their ID#, would be calculating their payments based upon their decisions, and that the experimenter in the room with them would never be able to associate their decisions with them (by ID# or otherwise), nor would he ever be informed of their individual payments.
6. A's were then read the instructions of the experiment, given a practice setting with which to familiarize themselves with the experiment, and were given several practice tasks to perform to ensure that they understood how to manipulate the division of tokens. This was the only time the experimenter would ever see their computer monitor.
7. A's recorded their decisions on paper "decision slips" which they then placed into manila envelopes. When all A's had completed all settings, the monitor carried these envelopes to the second experimenter in the second room. This experimenter then calculated the payments for group B first, placed them into security envelopes, handed them out and dismissed group B. He then calculated A's payments, placed them in security envelopes, and slid them under the door of the first room. The envelopes were then passed out to the A's who were then dismissed.
8. To further ensure experimenter blindness, subjects were informed that their data would not be recorded until all experimental sessions had been carried out.

These processes created a very transparent and trustworthy experimenter-blind set-up which eliminated experimenter effects as much as possible and ensured that, if dictators showed any appearance concerns, they did so in the interest of appearing a certain way to their partner, and not to the experimenter. Detailed instructions in the form of the scripts read by the experimenters to the subjects are available in Appendix A along with screenshots of the experimental interface.

B. Experimental Design

HMS and BF provide strong evidence that recognition affects behavior, but given their between subject design such a preference cannot be identified within a single subject. Furthermore, in the experiments utilized by HMS and BF, giving more is always – at least weakly – perceived as being more altruistic; as such, while the experiments of HMS and BF are very effective in identifying a recognition effect (as being present or not), they are unable to identify the strength of this effect.[23]

The goal of the experiment was to create the trade-off necessary to determine the relative importance each subject places on *creating* certain allocations relative to *being perceived as having created* certain allocations. This is done using the following three departures from the standard DG:

1. Incorporating a second, "anonymous" pile of tokens

A second sum of tokens is given to D such that she must now decide how to divide two separate amounts, by passing y from the first of size n, and x from the second of size m. The existence of this second sum of tokens is unknown to R, and any amount passed by D from this amount is given to R by the experimenter in the form of a "show-up payment." D, therefore, receives no recognition for any amount, x, passed from m. D is

[23] This is true of all DG studies.

aware of the manner in which R will be paid and of R's ignorance of the m additional tokens.

2. Using Multiple Settings

D is asked to make ten decisions instead of the usual one (she must choose an x and a y ten times for ten different budgets, where "different budgets" is explained in "Varying Budget Constraints" below). Each budget is referred to by a unique setting number "setting 1-10." The order in which the settings are presented to subjects was randomly determined before each experimental session to avoid any ordering affects.

3. Varying Budget Constraints

The amounts passed, x and y, are multiplied by positive numbers a and b (respectively) before being given to R. a and b are henceforth referred to as the "multipliers." These multipliers, as well as the magnitudes of m and n, vary between decisions, creating 10 different, overlapping, budget constraints. a and b represent the inverse of the prices of giving x and y respectively. m and n represent the income available for allocation.[24]

[Table 1 Here]

Table 2 lists the parameters defining the ten different budgets constraints. These budgets are numbered 1 through 10 for ease of reference, though the order they were presented to each dictator was altered between groups of subjects. [25]

[Table 2 Here]

[24] It should be noted that slide bars were used to determine the amount passed from each pile (x and y). Since the slide is always located somewhere along the bar, the starting location of the slide was randomly set for each dictator to avoid any status-quo bias.

[25] In addition to these 10 settings, dictators were given a practice setting to acquaint them with the experimental procedure. The practice settings asked subjects to "familiarize" themselves with the interface and to perform several tasks such as finding multiple ways of creating the same payoffs. This practice setting is provided in Appendix B.

The multipliers cover a wide range of possible values, including one very large value of a = 10 (Setting 9). The "very large" multipliers are used to elicit efficiency seeking behavior from subjects who may not have shown an interest in efficiency for smaller values of the multiplier.[26] Table 3 provides definitions for terminology which will be used throughout this paper. [27]

[Table 3 Here]

The term "true" (as in "true behavior) describes how the dictators' actions affect the *total* division of tokens. For example, dictators "behave truly selflessly" by setting $\pi_D < \pi_R$. The term "appears" describes how dictators' actions affect the distribution of tokens as *perceived* by the recipient. For example, dictators "appear selfless" by setting $\pi_D^K < \pi_R^K$.

IV. Results

A. Between Subjects Results

Figure I shows the distribution of total tokens earned by the recipient as a fraction of total earnings.

[Figure I Here]

In keeping with previous dictator game studies, dictators were primarily self-interested, electing to keep more than 50% for themselves in 197 of 280 allocations (70.3%), and electing to pass nothing in 32 allocations (11.4%). This is consistent with this group of subjects' behavior in the baseline dictator game in which 14.2% of offers were perfectly selfish (see Figure VII for the full distribution). A Wilcoxon-Mann-Whitney test of average percentage earned across all 10 settings by each recipient vs. the percentage earned by each recipient in the baseline dictator

[26] Subjects did, in fact, show the greatest deviation from otherwise consistent behavior in Setting 9.
[27] This paper uses the term "equitable" to describe payoffs which leave both D and R with equal earnings. It may be the case that equity is viewed by some as keeping and passing an equal number of tokens (not earning an equal amount of money). Subject behavior, however, indicates that this is probably not the case. For settings in which b≠1, only 4 decisions (1.7%) showed dictators dividing an equal number of tokens compared to 60 decisions (26.8%) in which dictators divided equal amounts of earnings.

game confirms that the distributions of total offers in the experimental settings and in the baseline dictator game are statistically indistinguishable ($z = -0.77$, $p = 0.4422$).[28] This data is also consistent with the findings of BF in their one-way identification and two-way identification settings which this study closely resembles in terms the level of social distance (i.e. interaction) between the dictator and the recipient prior to dictators selecting their divisions.[29]

To get a first look at the effect of recognition on altruistic behavior, Figure II presents the distribution of anonymous offers separately from open offers across all settings.

[Figure II Here]

The clear points of contrast in these two distributions are at the perfectly selfish (0%-2.5%) and egalitarian (47.5%-52.5%) divisions. The number of perfectly selfish anonymous offers is twice the number of perfectly selfish open offers. The number of equitable open offers is approximately seven times the number of equitable anonymous offers.

It is clear from these distributions that being openly egalitarian and anonymously selfish is, by and large, preferred to being anonymously egalitarian and openly selfish. However, since the anonymous and open offers are made simultaneously, it is likely that there is an interaction between the amount passed anonymously and the amount passed openly that we are ignoring by considering these offers separately. Furthermore, given the presence of multipliers, it is necessary that we also take into account possible efficiency motives which may be driving these results. For example, the number of anonymous offers greater than 52.5% is most likely due to a concern for efficiency (since eight of the ten settings incentivize anonymous giving using a/b >

[28] Similarly, a t-test of difference in means between percentage of total earnings received by the recipient in the experimental treatment (33.7) vs. percentage of total earnings received by the recipient in the baseline dictator game (32.7) returns an insignificant p-value of 0.8353.

[29] BF found 11% and 0% of offers to be perfectly selfish in their one-way identification and two-way identification settings respectively. More comparisons between this study and BF are available in Appendix B.

1). Indeed, of the 57 anonymous allocations leaving R with at least half of the total amount earned anonymously, 53 (92.9%) were made in settings where a>b and all but 1 were made when a>1. In contrast, in the baseline dictator game where all allocations maximize the sum of recipient and dictator earnings, no allocations greater than 50% were made. Interestingly, the number of open offers greater than 50%, appears to not have been created in the pursuit of efficiency. Of the 42 open allocations leaving R with at least half of the total amount earned openly, only 14 (33.3%) were made when b>1, and only 9 (21.4%) were made when b>a. To examine these "other behavioral motives" we need to think of the dictator's problem from a different perspective.

Given the informational structure of this experiment, the dictator's decision can be thought of as having to select a preferred point along a given budget constraint in the (π_D, π_R^A, π_R^K) plane. One such generalized budget constraint is illustrated in Figure III.

[Figure III Here]

Projecting this budget constraint onto the (π_R^A, π_R^K) plane gives us a space in which we can interpret our results from a behavioral perspective. In this space the origin represents a perfectly selfish allocation ($\pi_R = 0$), and the Northeast "corner" represents a perfectly selfless allocation ($\pi_D = 0$). We can partition this space into four quadrants with different behavioral interpretations by graphing the locus of allocations for which $\pi_D = \pi_R$ (the "True Equity Line") and the locus of allocations for which $\pi_D^K = \pi_R^K$ (the "Perceived Equity Line"). Figure IV presents a generalized budget constraint projected into the (π_R^A, π_R^K) plane with these loci illustrated. This results in the creation of four quadrants (indexed with Roman numerals), four line segments (indexed with Arabic numerals), and one point (indexed with the number 5).

[Figure IV Here]

Given this vantage-point we can think of any single allocation as being located either in one of the four quadrants (I, II, III, or IV),or along one of the four line segments (1, 2, 3, or 4) or at point 5. These quadrants and line segments are referred to hereafter as "areas of interest." Table 4 describes the true and perceived interpretations of allocations located within each area of interest and indexes the number of observations located in each area of interest.

[Table 4 Here]

Once again, we see a large number of selfish offers with 97 (34.6%) allocations located in Quad III. This is, however, a large departure from the previous examination of selfish behavior (Figure I) which showed 197 allocations (70.3%) as being selfishly motivated. The difference in these numbers comes from the behaviorally more accurate reinterpretation of the allocations as provided in Table 4. Of the 197 "truly selfish" allocations, 73 appeared equitable, 27 appeared selfless, and the remaining 97 appeared selfish.

Similarly, the 66 equitable offers (23.5%) shown in Figure I can be reinterpreted as being only partially motivated by an equitable appearance. Only 22 of these 66 offers appeared equitable to the recipient (Point 5), and – surprisingly – only 10 appeared selfless (Line 3). The remaining 34 (Line 4) actually appeared selfish to the recipient, showing that those who valued true equity did not usually concern themselves with their appearance.[30]

The tendency to desire a specific behavioral outcome (i.e. locating on a line segment as opposed to in a quadrant) is quite prevalent. If we consider the 32 perfectly selfish offers to be separate from other Quad III offers, the modal area of interest is along a line and not within a quadrant. Line 1 where the dictator appears egalitarian while still behaving selfishly, is the modal appearance/behavioral combination of choice, accounting for 26.1% of all allocations.

[30] Had they valued both true equity *and* a selfless appearance we would have observed more observations on line 3 than line 4, such allocations would have actually decreased π_R in the eight settings where a>b.

Combining the data along the lines and at point 5, 96 allocations (34.3%) appeared equitable while 66 (23.6%) were truly equitable. 22 (7.9%) both appeared equitable and were truly equitable, meaning that 140 (50%) allocations were either truly equitable, or equitable in appearance. This finding supports theories of fairness preferences, though apparently the prevalence of fair behavior is primarily driven by a desire to appear equitable, not to truly behave equitably.[31]

The above interpretations of the data provide support for the findings of previous dictator game studies, both in terms of the prevalence of "fair" offers and the importance of recognition. To build upon these findings we now present analysis which shows these influences to be present within subjects.

B. ***Within Subjects Results***

One way to measure consistency in this experiment is to examine how frequently a dictator chooses to locate inside a given area of interest. Figure V illustrates this consistency by tabulating each dictator's maximum number of allocations (out of 10) located in the same area of interest. Figure V shows that 64% of dictators locate in the same area of interest in 8 or more settings, and 80% of dictators locate in the same area of interest in at least five of the ten settings.

[Figure V Here]

To determine what behavioral motives are most influencing these allocations we can assign to each dictator allocation a vector with elements defined by the distances from the allocation to behaviorally salient points in the (π_R^A, π_R^K) plane. Let $A_{i,t,}$ represent the allocation

[31] The fairness model of Fehr and Schmidt (1999) can be easily adapted to accommodate such preferences by simply adding to their current model: $- \gamma * \Sigma \min\{ \pi_R^K - \pi_R^D\} - \delta * \Sigma \min\{ \pi_R^D - \pi_R^K\}$ to reflect a subject's distaste for perceived inequity in their favor as well as their distaste for perceived inequity in their partner's favor.

made by dictator i in setting t. For each $A_{i,t,}$ we assign four values: <$DS_{i,t}$, $DA_{i,t,}$, $DEq_{i,t,}$, $DEff_{i,t}$> which represent degree of self-interest (DS), degree of appearance-concern (DA), degree of true-equity-concern (DEq), and degree of efficiency-concern (DEff). $DS_{i,t}$ is defined as the distance from $A_{i,t,}$ to the origin of the (π_R^A, π_R^K) plane. $DA_{i,t}$ is the vertical distance from $A_{i,t,}$ to the "appearance of equity" line in setting t. $DEq_{i,t}$ is the distance between $A_{i,t,}$ and the nearest point on the "true equity" line in setting t. $DEff_{i,t}$ is the distance between $A_{i,t,}$ and the *non-axis* efficient border of the budget constraint in the (π_R^A, π_R^K) plane in setting t (i.e. the Eastern border when a>b and the Northern border when b>a). [32]

By calculating the average value of each component of this vector for each subject, we can identify which subjects are most concerned with certain types of outcomes. Table 5 gives these average vector values for each subject. Note that smaller values imply a greater importance since the values measure the distance between the allocation and the behaviorally salient point.

Using the average vector values for each subject we assign to each dictator a behavioral "type" corresponding to which element of their behavioral vector has the minimum average value. Each dictator is assigned exactly one of the following five types: "Self-interested," "Appearance-concerned," "True-Egalitarian," "Efficiency-concerned," "Non-Conformist."

Given that average values are used for this calculation, even a dictator who makes all decisions randomly will be assigned a type. To control for this, only dictators who focus five or more of their allocations in a single areas of interest are assigned a type other than non-conformist. Further, they will only be assigned a type other than non-conformist if the area of

[32] Note that the distance DEff does not measure a concern for pareto efficiency, but rather a desire to allocate earnings in the most productive way. This distinction is necessary to separate the perfectly selfish D's (all of whose allocations are pareto efficient) from D's whose decisions were most strongly based upon the value of the multipliers a and b.

interest in which they locate the majority of their allocations is consistent with the minimum element of their average behavioral vector. To be clear, dictators who fail to focus the majority of their allocations in an area of interest consistent with the minimum valued element of their average behavioral vector are labeled "non-conformists." Table V presents the "type" assigned to each subject. Figure VI(A) provides sample dictator behavior across various budget constraints for the purpose of illustration. Figure VI(B) provides all dictator behavior across all budget constraints.

[Table 5 Here]

[Figures VI(A) and VI(B) Here]

Given this method of determining types it was always possible to distinguish between the self-interested, the appearance concerned, and the true egalitarians, but it was not always possible to distinguish between these three types and the efficiency concerned. The four subjects who stood out as "efficiency concerned" (subjects 1, 4, 13, and 20 – see Table 5) were all also strongly of the type "true egalitarian." As such, these four subjects were labeled "true egalitarians" for the sake of between type analysis. Nine subjects (32.1%) were of type "self-interested," nine (32.1%) were of type "appearance-concerned," six (21.4%) were of type "true-egalitarian," and four (14.3%) were "non-conformists."[33]

C. Between Group (Types) Results

An examination of average behavior by group in terms of the amount passed anonymously, openly, in total, and in the baseline dictator game is presented in Table 6.

[33] These types explained the majority of allocations in each area of interest. 82% of the 66 allocations on Lines 3, 4 or point 5 were made by the true egalitarians. 75% of the 100 allocations on line 1 or Quad II were made by the appearance-concerned, and 80% of the 97 allocations in Quad III were made by the self-interested. Interestingly, 80% of the Quad I offers (4 of 5), 36.4% of the Quad IV offers (4 of 11), and 37% of the Quad II (10 of 27) offers as well as the only Line 2 allocation were made by non-conformists. Non-conformists had an average H of 0.25.

[Table 6 Here]

It is impossible to statistically distinguish between the behavior of the appearance concerned and the self-interested when observing only the average amount passed anonymously by these two groups (13.8% and 11.7% respectively, t=0.212). However, the average amount passed openly by the appearance concerned is statistically much greater than the average amount passed openly by the self-interested (48.7% and 20.4% respectively, t=3.641). This desire to be recognized as more giving is enough to create a statistically significant difference in the average amount passed in total between these two groups (35.1% and 15.6%, t=3.211).

The average amount passed anonymously can be used to distinguish between the true-egalitarians and the appearance concerned (67.5% and 13.8% respectively, t=4.313). Interestingly, the true-egalitarians pass significantly less than the appearance concerned openly (29.2% vs. 48.7%, t=1.890). This dichotomous result is due to the true egalitarians valuing efficiency but not appearance. True egalitarians, on average, pass a significantly larger share of the total amount relative to the appearance concerned (51.5% vs. 35.1%, t=3.796).

When comparing the behavior of true egalitarians with the self-interested, the average amount passed openly by these two groups is statistically indistinguishable (29.2% and 20.4%, t=0.739), while the average amount passed anonymously by the true egalitarians is much greater (67.5% vs. 11.7%, t=4.467). The average amount passed in total by the true egalitarians is, of course, statistically greater than the average total amount passed by the self-interested (51.5% vs. 15.6%, t=5.776).

Finally, comparing average behavior by group within the setting of the standard dictator game, where all dictator behavior is observed by the recipient, the average behavior of true-

egalitarians and the appearance concerned are no longer statistically distinguishable (mean total offers are 49.8% and 42.2% respectively, t=1.264). Both these groups, however, pass statistically more than the self-interested, who share only 12.3% in the baseline dictator game.

V. Conclusion

When failing to account for the role of recognition on altruism, baseline dictator games are inclined to overstate the importance of a preference for equality. What is commonly observed as a "fairness preference" is more often than not being driven by a preference for appearing fair. None-the-less, a significant share of subjects does, in fact, value true equity in the sense of equal total earnings.

By openly passing shares of wealth of approximately 50%, the appearance concerned subjects tacitly admit their knowledge of a 50-50 norm in giving. Whether these subjects actually value true equality cannot be determined, though it is clear that they value being perceived as one who values equality.

The self-interested appear unaffected by the social stigma the appearance-concerned are trying to avoid by passing large shares openly. Whether this is due to obliviousness of social norms or simply an overwhelming desire to maximize personal wealth is unclear from this research. Future research may be interested in determining the extent to which this group is immune to the impressions they make on others and whether such subjects will behave with less regard to the concerns of others in more strategic environments. Other research may endeavor to allow subjects to costlessly adjust their appearance in the hopes of determining whether the selfish are unaware or just highly self-interested.[34]

[34] This could be done by asking subjects to play a standard dictator game, then – after allocations have been decided – giving subjects the opportunity to change the amount that the recipient will be told was available for division. This experiment would also allow the researcher to determine whether the 50-50

Given the small sample size of this experiment, the distribution of behavioral types across the population cannot be reliably estimated. Furthermore, were more subjects sampled, other behavioral patterns may have emerged. Interesting future research may be found in simply replicating this experiment with a larger sample. It would also be interesting to see how behavior changes with differing recipient characteristics, particularly if the recipient is a well known charity (i.e. Eckel and Grossman (1996)). Also, examining whether recognition effects are prevalent in the field would be fruitful as it would help to determine whether the appearance concerned consciously sort themselves out of voluntary giving opportunities (since, despite the term "voluntary" not all such scenarios can be avoided, i.e. the tipping scenario at the opening of this paper). A field experiment examining the affect of recognition on "voluntary" giving in avoidable and less-avoidable settings would therefore be very interesting.

I would like to acknowledge the Russell Sage Foundation for their support which made this research possible. Without their small grant for behavioral economics, there is no way that I would have been able to conduct this experiment.

---

norm is arrived at out of preferences for equality or out of preferences for appearing as generous as possible without sacrificing "too much" of the wealth.

CHAPTER 2 - BIBLIOGRAPHY

Anderson, Rodgers, and Rodriguez; "Cultural Differences in Attitudes Toward Bargaining," *Economics Letters* 69, 2000, pp. 45-54

Andreoni, J.; "Impure Altruism and Donations to Public Goods: A Theory of Warm Glow Giving," *Economic Journal*, 1990, pgs. 464-477

Andreoni, J., Bernheim, B.; "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects," *Working Paper*, April 2006

Andreoni, J., and Miller, J.; "Giving According to GARP," *Econometrica*, Vol. 70, No. 2, March 2002, pgs. 737-753

Andreoni, J., and Vesterlund, L.; "Which is the Fair Sex? Gender Differences in Altruism," *The Quarterly Journal of Economics*, February 2001

Batson, C., Thompson, E., and Chen, H.; "Moral Hypocrisy: Addressing Some Alternatives," *Journal of Personality and Social Psychology*, Vol. 83, No. 2, 2002, pgs. 330-339

Batson, C., Thompson, E., Seuferling, G., Whitney, H., and Strongman, J.; "Moral Hypocrisy: Appearing Moral to Oneself Without Being So," *Journal of Personality and Social Psychology*, Vol. 77, 1999, pgs. 525-537

Becker, G.; "A Theory of Social Interactions," *Journal of Political Economy*, vol. 82, 1974, pgs. 1063-1093

Binmore, K., Shaked, A., and Sutton, J.; "Fairness or Gamesmanship in Bargaining: An Experimental Study," ICERD Discussion Paper, London School of Economics, 1984

Burnham, T.; "Engineering Altruism: A Theoretical and Experimental Investigation of Anonymity and Gift Giving," *Journal of Economic Behavior and Organization*, vol. 50, 2003, pgs. 1-13

Cameron, Lisa; "Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia," *Economic Inquiry*, Vol. 37, No. 1, Jan. 1999, pp 47-59

Charness, G., and Gneezy, U.; "What's in a name? Anonymity and Social Distance in Dictator and Ultimatum Games," *Working Paper*, August 2003

Clotfelter, C.; "Federal Tax Policy and Charitable Giving," *Chicago: University of Chicago Press*, 1985

Dana, J., Cain, D., and Dawes, R.; "What You Don't Know Won't Hurt Me: Costly (but quiet) exit in DGs," *Carnegie Mellon University*, October 12, 2005 version

Eckel, C. and Grossman, P.; "Altruism in Anonymous DGs," *Games and Economic Behavior*, Vol. 16, 1996, pgs. 181-191

Dickinson, D., Tiefenthaler, J.; "What is Fair? Experimental Evidence," *Southern Economic Journal*, Vol. 69, No. 2, Oct. 2002, pgs. 414-428

Eckel, C., and Grossman, P.; "Chivalry and Solidarity in Ultimatum Games," *Working Paper*, Virginia Polytechnic Institute and State University, 1997

Eckel, C., and Grossman, P.; "Are Women Less Selfish than Men? Evidence from Dictator Experiments," *Economic Journal*, CVII, 1997, pgs. 726-735

Fehr, E., and Schmidt, K.; "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics*, vol. 114, 1999, pgs. 817-868

Forsythe, R., Horowitz, J., Savin, N., Sefton, M.; "Fairness in Simple Bargaining Experiments," *Games and Economic Behavior*, Vol. 6, July 1994, pgs. 347-369

Geanakoplos, J., Pearce, D., and Stacchetti, E.; "Psychological Games and Economic Rationality," *Games and Economic Behavior*, Vol. 1, 1989, pgs. 60-79

Guth, Werner, Schmittberger, and Schwartz; "An Experimental Analysis of Ultimatum Bargaining," *Journal of Games and Economic Behavior*, December 1982, pp. 367-388

Henrich, J.; "Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon," *The American Economic Review*, Vol. 90, No. 4, Sep. 2000, pp. 973-979

Henrich, Boyd, Bowles, Camerer, Fehr, Gintis, and McElreath; "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies," *AEA Papers and Proceedings*, May 2001

Hoffman, E., McCabe, K., Shachat, K., and Smith, V.; "Preference, Property Rights, and Anonymity in Bargaining Games," *American Economic Review*, vol. 7, 1994, pgs. 346-380

Hoffman, E., McCabe, K., and Smith, V.; "Social Distance and Other Regarding Behavior in DGs," *American Economic Review*, vol. 86, 1996, pgs. 653-660

Kahneman, D., Knetsch, J., and Thaler, R.; "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *The American Economic Review*, Vol. 76, No. 4, 1986, pgs. 728-741

Krupka, E., and Weber, R.; "The Focusing and Observational Effects of Norms," *December 2005*

Konow, J.; "A Positive Theory of Economic Fairness," *Journal of Economic Behavior and Organization*, Vol. 31, 1996, pgs. 13-35

Konow, J.; “Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions,” *The American Economic Review*, Vol. 90, No. 4, Sep. 2000, pgs. 1072-1091

Lazear, E., Malmendier, U., and Weber, R.; “Sorting in Experiments with Application to Social Preference,” *Working Paper*, June 2005

Munier, B. and Zaharia, C.; “High Stakes and Acceptance Behavior in Ultimatum Bargaining: A Contribution from an International Experiment,” *Theory and Decision*, 53, 2003, pp. 187-207

Rabin, M.; Rabin, M.; “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, vol. 83, 1993, pgs. 1281-1302

Roberts, R.; “A Positive Model of Private Charity and Public Transfers,” *Journal of Political Economy,* Vol. 92, 1984, pgs. 136-148

Rosenblat, T. “Why Good Looks Matter,” *mimeo.* 2001

Roth, and Erev; “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term,” *Games and Economic Behavior*, 8, 164-212

Roth, Prasnikar, Okuno-Fujiwara, and Zamir; “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” *The American Economic Review*, Vol. 81, No. 5, Dec. 1991, pp 1068-1095

Roth, and Slonim; “Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic,” *Econometrica*, Vol. 66, No. 3, May 1998, pp 569-596

Solnick, S.; “Gender Differences in the Ultimatum Game,” *Working Paper*, University of Miami, 1997

Solnick, S. and Schweitzer, M.; “The Influence of Physical Attractiveness and Gender on Ultimatum Game Decisions,” *Organizational Behavior and Human Decision Processes*, vol. 79, 1999, pgs. 199-215

Stahl, I.; “Bargaining Theory,” Stockholm: Economic Research Institute, 1972

CHAPTER 2 – TABLES

Table 1: Notation

Expression	Definition
m	Number of tokens available in Left (Anonymous) Pile
n	Number of tokens available in Right (Open) Pile
x	Amount passed from Left (Anonymous) Pile
y	Amount passed from Right (Open) Pile
a	Amount by which tokens passed from Anonymous Pile are multiplied
b	Amount by which tokens passed from Open Pile are multiplied
$\pi_D = (m-x) + (n-y)$	Total Dictator Earnings
$\pi_D^K = (n-y)$	Open Dictator Earnings
$\pi_D^A = (m-x)$	Anonymous Dictator Earnings
$\pi_R = a*x + b*y$	Total Recipient Earnings
$\pi_R^A = a*x$	Anonymous Recipient Earnings (received as a "show up" payment)
$\pi_R^K = b*y$	Open Recipient Earnings

Table 2: Dictator Budgets (Settings)

Setting #	m	n	a	b	a/b
1	100	210	2	1	2
2	40	200	5	1	5
3	170	100	1.5	0.5	3
4	130	130	2	1.5	1.33
5	60	200	2.5	1	2.5
6	220	100	1	2	0.5
7	100	200	1.5	1	1.5
8	120	140	3	0.5	6
9	40	140	10	1	10
10	160	140	1.5	2	0.75

Table 3: Summary of Behavioral Terminology

Term	Definition
Truly Selfless	$\pi_D < \pi_R$
Truly Equitable	$\pi_D = \pi_R$
Truly Selfish	$\pi_D > \pi_R$
Appears Selfless	$\pi_D^K < \pi_R^K$
Appears Equitable	$\pi_D^K = \pi_R^K$
Appears Selfish	$\pi_D^K > \pi_R^K$

Table 4: Dictator Behavior Given an Allocation in (π_R^A, π_R^K) Plane (Corresponds with Figure VI)

Area of Interest	Appearance	True Behavior	# (%) of Allocations (N = 280)
Quad I	Selfless	Selfless	5 (1.8%)
Quad II	Selfless	Selfish	27 (9.6%)
Quad III	Selfish	Selfish	97 (34.6%)
Quad IV	Selfish	Selfless	11 (3.9%)
Line 1	Equitable	Selfish	73 (26.1%)
Line 2	Equitable	Selfless	1 (0.3%)
Line 3	Selfless	Equitable	10 (3.6%)
Line 4	Selfish	Equitable	34 (12.1%)
Point 5	Equitable	Equitable	22 (7.9%)
Origin ($\pi_R = 0$)*	Selfish	Selfish	32 (11.4%)

*These observations are double counted as they also fall into Quad III

"Lines" are given a margin of error of +/- 2.5% (so an allocation between 47.5% and 52.5% would be considered equitable). This is done for two reasons. Firstly, given the multipliers it was not always possible to set earnings (true or perceived) exactly equal. Secondly, many subjects (71.4%) showed a preference for receiving an "even amount" (an amount ending in $X.50 or $X.00) in 6 or more settings, and I did not want a distaste for dimes to be confused with a distaste for equity.

Table 5: Subject Types, Area-of-Interest Frequencies, and Average Behavioral Vector Values

Subject #	Type	Mean DS	Mean DA	Mean DEq	Mean DEff	Majority (count)
1	True Egalitarian	204.3512	69.35	32.9	25.69827	Line 4 (6)
2	True Egalitarian	112.8029	14.8	2.25	113.3255	Point 5 (10)
3	App. Concerned	80.15	14.7	113.65	164.0255	Line 1 (10)
4	True Egalitarian	155.117	63.15	10.6	59.87554	Line 4 (9)
5	Non-conformist	110.0636	30.25	34.35	125.5755	None
6	App. Concerned	124.3	30.45	129.8	151.2983	Line 1 (5)
7	Self-Interested	10	80.75	256	170.8255	Quad III (9)
8	Non-conformist	103.896	29.65	51.9	150.0255	None
9	Self-Interested	61.6	31.15	148.8	168.3255	Quad III (9)
10	True Egalitarian	154.7914	31.8	26.6	106.8483	Line 5 (8)
11	App. Concerned	90.42852	28.45	77.75	150.5755	Line 1 (5)
12	Self-Interested	88.1379	42.85	102.7	108.5755	Quad III (8)
13	True Egalitarian	175.2393	64.05	0.8	9.168398	Line 4 (7)
14	Self-Interested	38.69093	72.15	209	146.4755	Quad III (9)
15	Non-conformist	110.6076	52.6	74.15	98.02554	None
16	Non-conformist	124.6849	30.35	49.5	113.7755	None
17	Self-Interested	63.68691	47.05	131.65	146.4755	Quad III (6)
18	Self-Interested	61.91577	49.45	146.2	162.5255	Quad III (7)
19	Self-Interested	0	90.75	270	180.8255	Quad III (10)
20	True Egalitarian	206.895	71.5	29.25	25.9816	Line 4 (6)
21	App. Concerned	90.66173	21.85	70.75	154.1755	Quad II (5)
22	App. Concerned	77.42444	21.6	111.85	162.2255	Line 1 (8)
23	App. Concerned	92.58981	18	64.25	137.4255	Line 1 (5)
24	Self-Interested	58.38516	33.75	154.8	165.3255	Quad III (10)
25	App. Concerned	122.7906	18.4	109.45	153.7983	Line 1 (8)
26	App. Concerned	95.64239	14.85	51.15	136.2755	Line 1 (9)
27	Self-Interested	0	90.75	270	180.8255	Quad III (10)
28	App. Concerned	80	14.75	114	163.8255	Line 1 (10)

A verbal description of each subject's behavior is provided in Appendix A

Table 6: Group Averages (columns 1-3) and Baseline DG Behavior by Group

	Group Averages	$\pi_R^A/(\pi_R^A + \pi_D^A)$	$\pi_R^K/(\pi_R^K + \pi_D^K)$	$\pi_R/(\pi_R + \pi_D)$	Baseline DG
A)	Average of Appearance Concerned (n = 9)	13.8 (22.2)	48.7 (13.7)	35.1 (10.3)	42.2 (15.6)
B)	Average of True Egalitarians (n = 6)	67.5 (29.7)	29.2 (29.1)	51.5 (5.5)	49.8 (0.17)
C)	Average of Self Interested (n = 9)	11.7 (22.3)	20.4 (20.6)	15.6 (15.5)	12.3 (13.5)
	t-stat (A vs. B)	-4.313***	1.890*	-3.796***	-1.264
	t-stat (A vs. C)	0.212	3.641***	3.211***	4.553***
	t-stat (B vs. C)	4.467***	0.739	5.776***	7.033***

Significant for two-tailed t-test for *α=0.10 **α=0.05 ***α<0.01

CHAPTER 2 – FIGURES

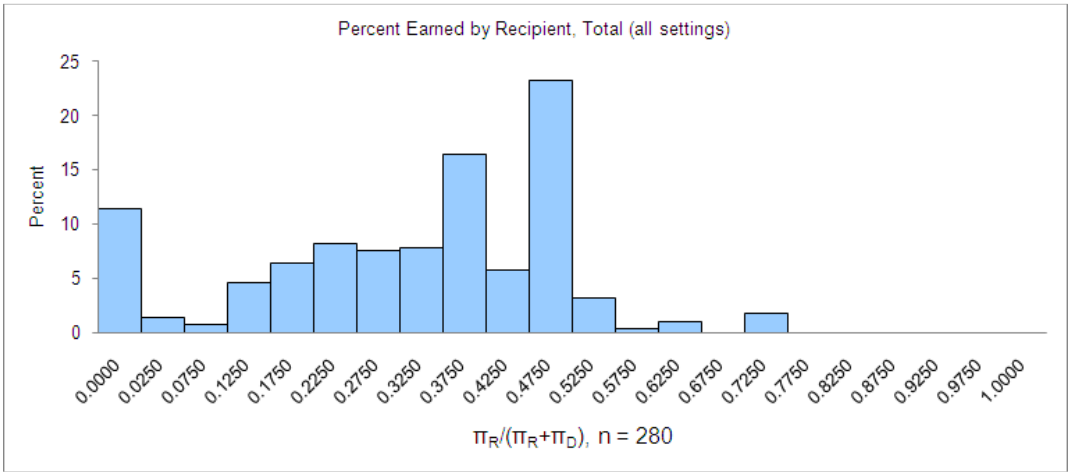

Figure 1: Percent Earned by Recipient: Total, All Budgets

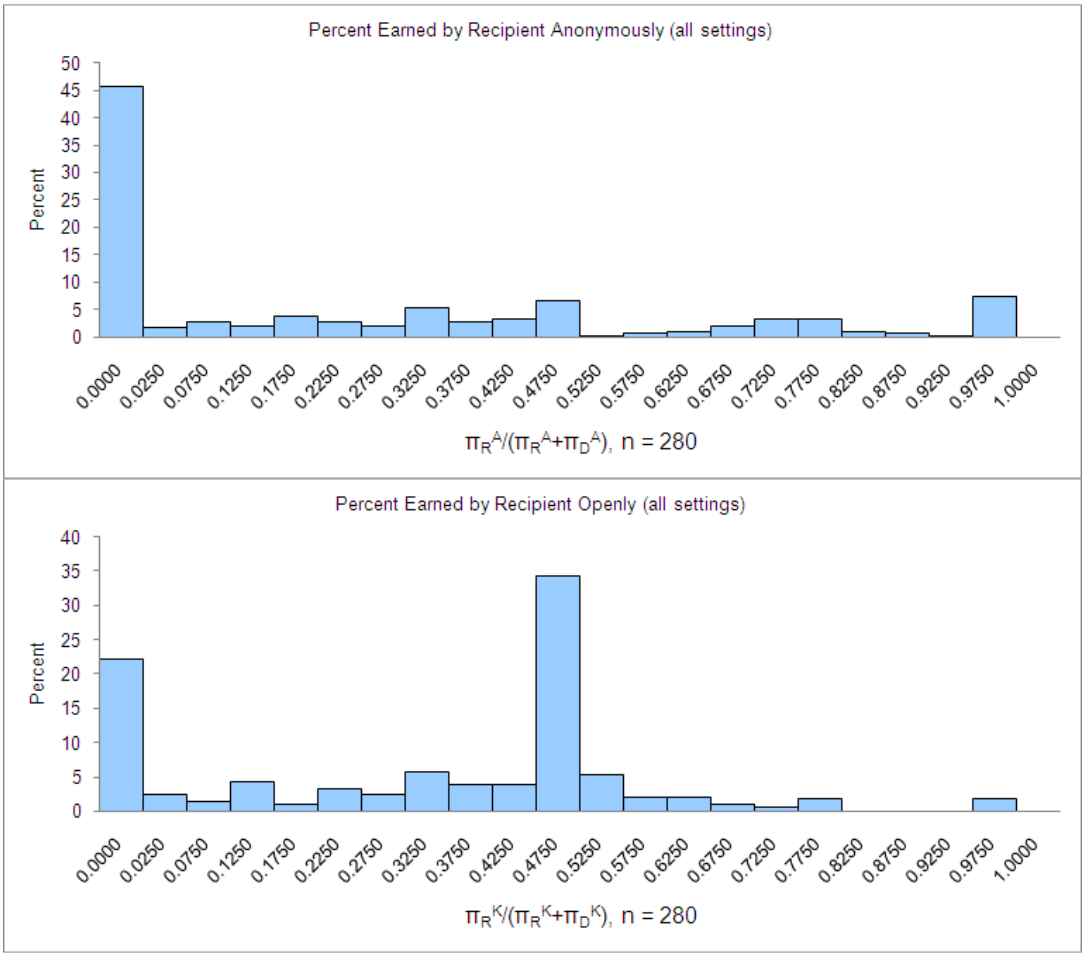

Figure 2: Percent Earned by Recipient: Anonymously (Above), Openly (Below), All Budgets

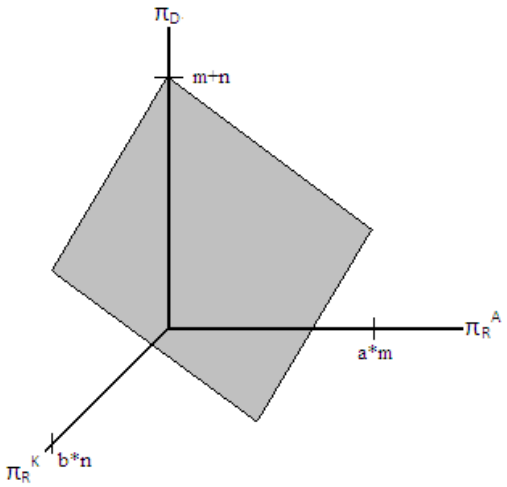

Figure 3: Generalized Dictator Budget Constraint

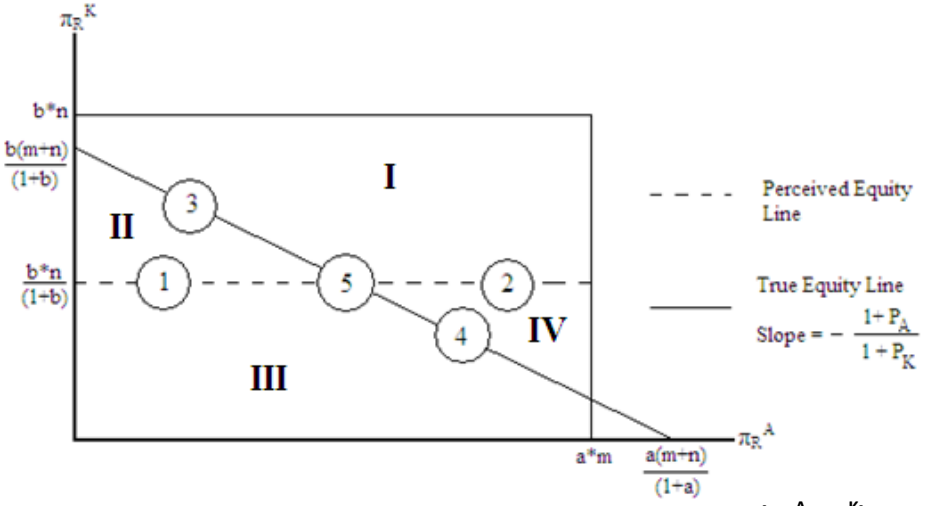

Figure 4: Generalized Dictator Budget Constraint Projected onto (π_R^A, π_R^K) Plane

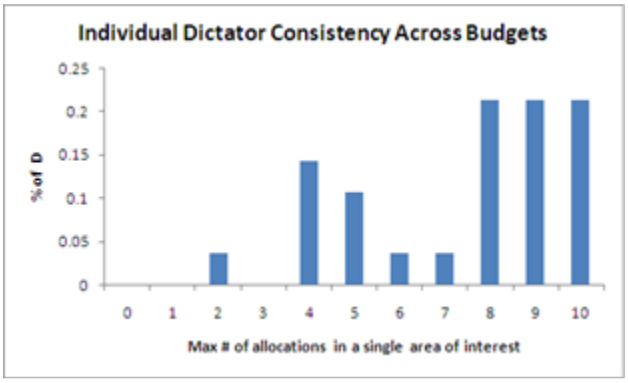

Figure 5: Dictator Consistency Across All Budgets

Given there are 9 areas of interest and 10 budgets, the minimum maximum number of allocations which could be located in a single area of interest is 2. A subject randomly allocating wealth would on average locate 2.87 of their allocations in the same area of interest (SD = 0.740). Subjects in this experiment, on average, located 7.43 of their allocations in the same area of interest.

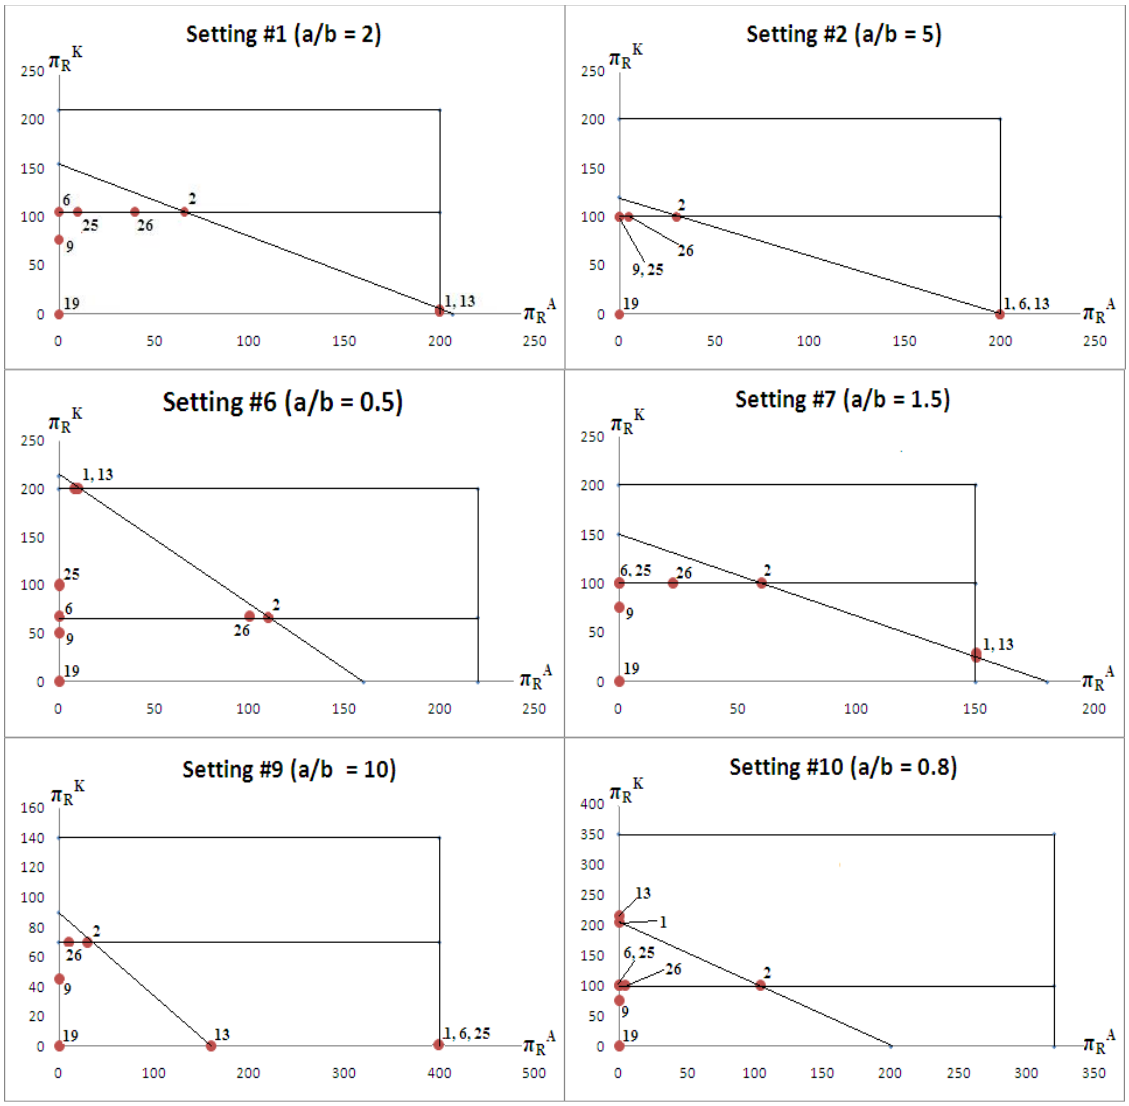

Figure 6A - Representative Subject Behavior Over Representative Budget Constraints*

*Subject ID's are included to note consistency across budgets. These eight subjects were selected because they provide clear illustrations of certain types of behavior. As shown in Figure V and Table 5, most subjects behaved quite consistently and could have been used to provide equally illustrative behavior.[35] Eight subjects were selected for the sake of simplicity and to avoid the figures from being messy or confusing.[36] Figure VI(B) provides all subject behavior over all budgets, but removes labels.

[35] Appendix A contains descriptions of all subject behavior.

[36] Description of Figure VI and other subject behavior: Here we see that Subject 19 is "perfectly selfish," passing nothing anonymously or openly in any setting. Subject 9 is "weakly selfish" and appearance concerned, never passing any amount anonymously, and always passing an amount openly that creates an appearance of passing an average of 39% across all settings, when in reality she is passing on average only 23%. Subject 13 is a true egalitarian who is also concerned with efficiency. This subject sets $\pi_R = \pi_D$ in every setting, and does so in the most efficient way by passing more anonymously when a > b and more openly when b > a. It is also important to note that they do so by passing as little as possible from whichever pile has the smaller multiplier (their setting #7 allocations illustrates this well). Subject 1 is very similar to Subject 13 (choosing to be truly equal in all settings and to do so in an efficient manner) except in setting 9. Here, Subject 1 passes 2.4 times as much as necessary to be truly equal. This shows that

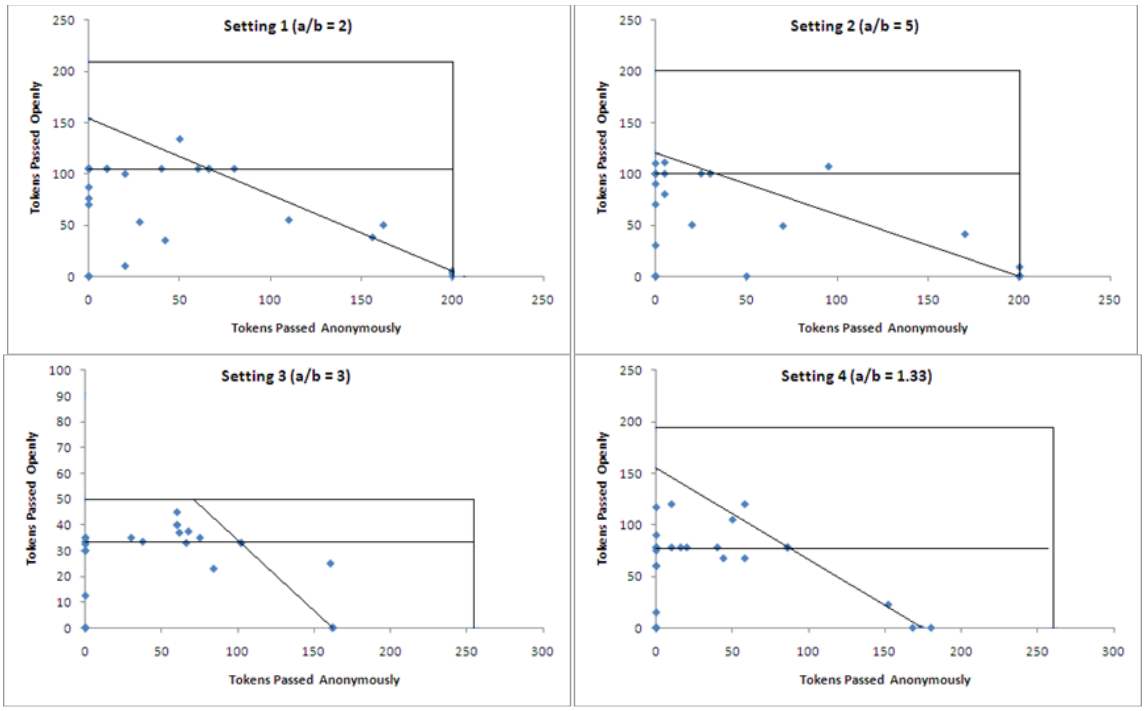

Figure 6B: All Subject Behavior Over All Budgets

Subject 1 values both equality and efficiency (like Subject 13), but will forego true equality in the name of efficiency if the incentives are strong enough (in this case, a multiplier of 10).

Subject 2 is not at all concerned with efficiency, but is concerned with both appearing egalitarian and truly behaving in an egalitarian manner. Likewise, subject 26 is not at all concerned with efficiency, but is concerned with appearing egalitarian, passing $\pi_R^K = \pi_D^K$ in all settings. Strangely, Subject 26 always passes some amount anonymously, but seems to prefer passing inefficiently, passing very little anonymously when a/b=10, but quite a bit when a/b = 0.5. This may be due to a "relativist" concern, i.e. wanting the recipient to always earn some fraction of what they earn. Indeed, upon a closer examination of this subject's data it is clear that they set recipient earnings to approximately 42% of their earnings in every setting except for setting #10.

Subject 25 values an egalitarian appearance as well as efficiency. This subject sets $\pi_R^K = \pi_D^K$ in most settings, but makes exceptions when a/b is very small or very large, passing more than half openly when b=2 and a=1 (setting #6) and nothing openly – but everything anonymously – when b=1 and a=10. Subject 6 is quite similar, setting $\pi_R^K = \pi_D^K$ in most settings, except when efficiency motives are strong enough. Subject 6 seems to require a multiplier of 5 or more to be persuaded to set $\pi_R^K \neq \pi_D^K$. In settings 2 and 9 (a=5, a=10 respectively) subject 6 foregoes their appearance concerns opting instead to pass 100% of their anonymous tokens.

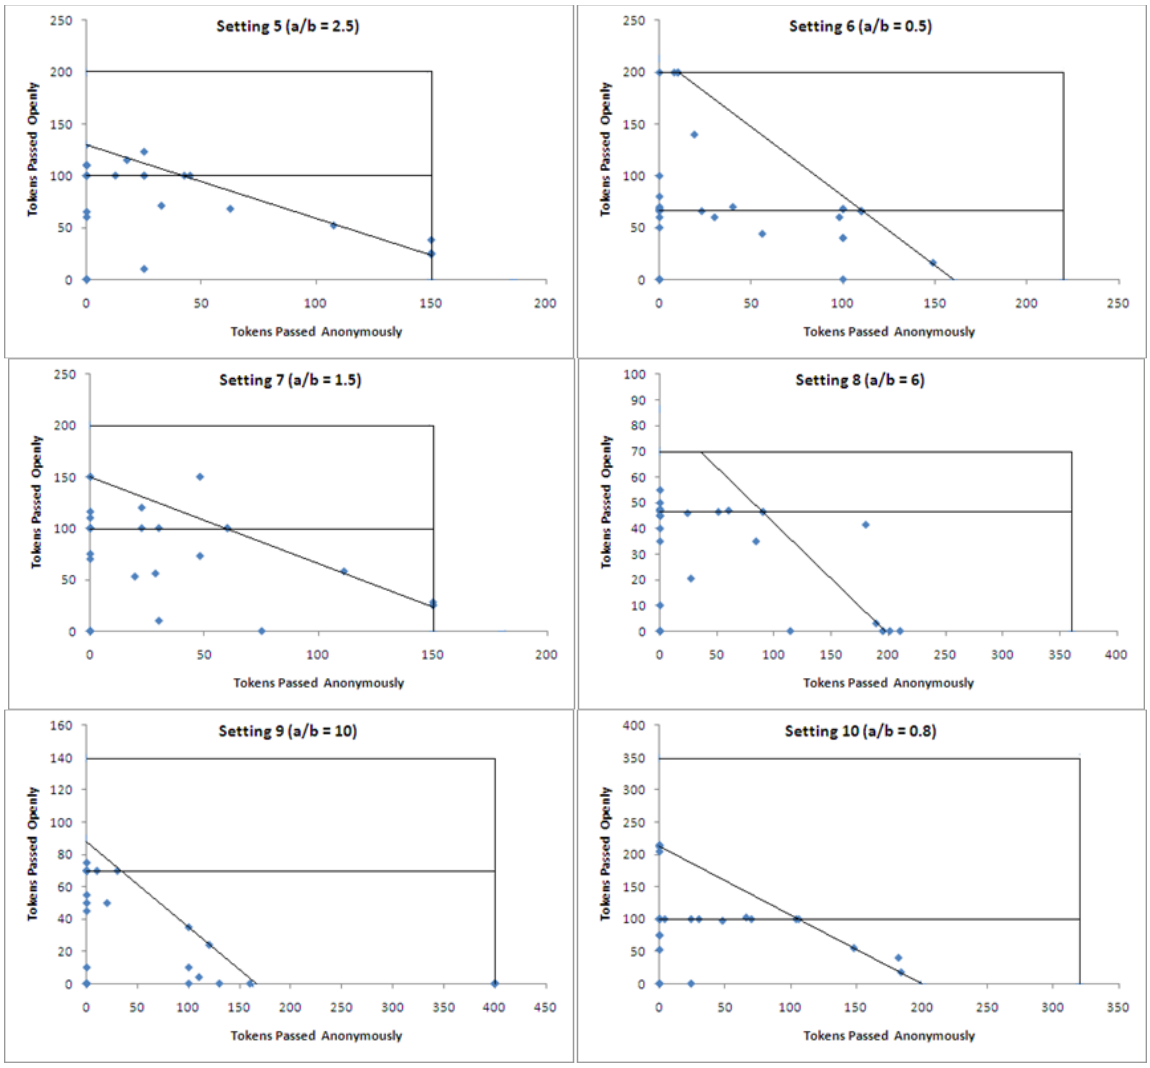

Figure 6B: Continued

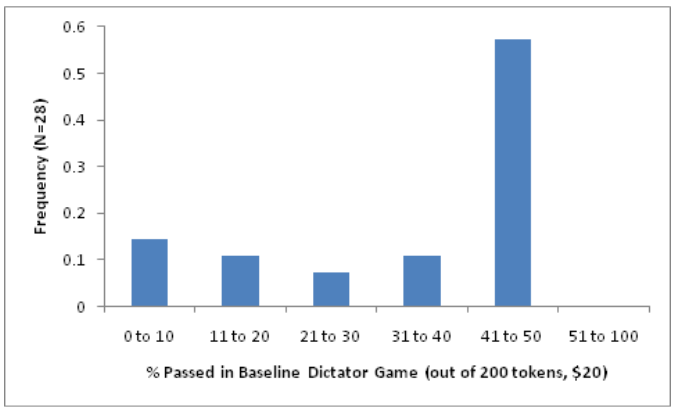

Figure 7: Baseline Dictator Game Results

Average percentage passed = 32.9% ($6.58 out of $20)

CHAPTER 2 – APPENDIX A – SPECIFIC SUBJECT BEHAVIOR DISCRIPTIONS

True Egalitarians

Subject #1 – Located on Line 4 seven times, when a > b. Located on Line 3 two times when b > a. The remaining time was in setting 9 (where a = 10) when he located in Quad IV. This subject clearly values efficiency and egalitarianism and shows no concern for appearances.

Subject #2 – Located at point 5 in all ten settings. Values being egalitarian *and* appearing egalitarian.

Subject #4 – Locates on Line 4 in nine of ten settings. Values being egalitarian.

Subject #10 – Locates at point 5 in 8 of ten settings. The remaining two are efficiency concerns (setting 2 and setting 9 when a = 5 and 10 respectively).

Subject #13 – Located on Lines 3, 4, or point 5 in all settings, showing a preference for equal true divisions. Located on Line 4 in seven settings where a > b, Line 3 in the 2 settings where b > a, and Line 5 in setting 4 where a/b = 1.33 (the closest a/b gets to 1 with a > b), showing efficiency concerns.

Subject #20 – Locates on Line 4 in seven settings where a > b, on Line 3 in the two settings where b > a, and in Quad IV in setting 9 where a = 10. Shows strong concern for egalitarianism and efficiency, and no concern for appearances.

Appearance Concerned

Subject #3 – Locates on Line 1 in all ten settings, and always passes nothing anonymously. Values appearing egalitarian, not being egalitarian, not efficiency.

Subject #6 – Locates on Line 1 in seven of ten settings. Makes exceptions in the settings where a > 4 by locating on Line 4 in setting 2 (a = 5) and Quad IV in setting 9 (a = 10). For low multipliers, this subject values appearing egalitarian, but shows a concern for efficiency as the multiplier

increases. At first, when a = 5, they keep total earnings equal but sacrifice appearances . Then, when a = 10, they given even more, allowing the recipient to earn more than they do, sacrificing appearances even more.

Subject #11 – Located on Line 1 in six of ten settings, and in Quad II in two settings, showing a concern for a selfless/egalitarian appearance. Located in Quad III in setting 6 (a = 1, this is the lowest value of a in all settings). Located on Line 4 in setting 1.

Subject #21 – Locates on Lines I and III or Quad II in nine of ten settings, showing a concern for an egalitarian/selfless appearance. Line 1 in five settings, Line 3 in two settings, and Quad II in two settings.

Subject #22 – Locates on Line 1 in eight of ten settings showing a concern for an egalitarian appearance. Shows no preference for efficiency or true equity.

Subject #23 – Passes an average of 49.2% openly across all settings, 38.7% total.

Subject #25 – Locates on Line 1 in eight of ten settings showing a preference for an egalitarian appearance with selfish motives. Locates in Quad IV in setting 9 where a = 10, and in Quad II where the ratio a/b = 1/2 (it's smallest value), showing a concern for efficiency at the extremes.

Subject #26 – Locates on Line 1 in eight of ten settings. Locates on Point 5 in two of ten, and therefore appears equitable in all settings.

Subject #28 – Locates on Line 1 in all ten settings, never passes anonymously.

<u>Efficiency Concerned</u>

Subject #1 – Locates on Line 4 seven times, when a > b. Located on Line 3 two times when b > a. The remaining time was in setting 9 (where a = 10) when he located in Quad IV. This subject clearly values efficiency and egalitarianism and shows no concern for appearances.

Subject #10 – Locates at Point 5 in eight of ten settings, making exceptions for setting 2 and setting 9, in which they give only anonymously, showing a concern for efficiency in terms of a but not a/b.

Subject #13 – Locates on Lines III, IV, or V in all settings, showing a preference for equal true divisions. Located on Line 4 in seven settings where a > b, Line 3 in the 2 settings where b > a, and Point 5 in setting 4 where a/b = 1.33 (the closest a/b gets to 1 with a > b).

Subject #20 – Locates on Line 4 in seven settings where a > b, on Line 3 in the two settings where b > a, and in Quad IV in setting 9 where a = 10. Shows strong concern for egalitarianism and efficiency, and no concern for appearances.

Selfish

Subject #7 – This was the archetypical "perfectly selfish" subject. They therefore located in Quad III in all settings, specifically, at the origin.

Subject #9 – Appearance concerned with selfish motives. Located in Quad III in all ten settings, did so by passing zero anonymously and roughly 30% openly in all settings.

Subject #12 – Located in Quad III in all settings except setting 6 in which a/b = 1/2, here this subject located in Quad II, showing that given the proper multiplier on open giving, this subject would elect to appear selfless(but would remain selfish).

Subject #14 – Locates in Quad III in nine of ten settings by setting recipient earnings to approximately 14% of total earnings. Locates on Line 1 in setting 6 where a/b = 1/2 (the smallest value this ratio takes in all settings).

Subject #17 – Locates in Quad III in 6 settings, passing an average of 24% across all settings. Locates in Quad II in one setting and Line 1 the remaining 3.

Subject #18 – Locates in Quad III in 9 settings, passing an average of 25% of the total across all settings. Locates in Quad I in setting 2, offering 81.2% of the anonymous amount and 53.5% of the open amount.

Subject #19 – Another "perfectly selfish" subject, locating on the origin in all ten settings.

Subject #24 – Locates in Quad III in all ten settings. Does so by never passing anonymously and passing about 25% of total earnings openly in all settings.

Subject #27 - Another "perfectly selfish" subject, locating on the origin in all ten settings.

Non-Conformists

Subject #5, Subject #8, Subject #15, Subject #16

CHAPTER 2 – APPENDIX B – EXPERIMENTAL MATERIALS

I. Instructions to subjects –*These were read to both groups of subjects at the start of the experiment*

Hello and thank you for participating in this study. Before we get going, I would like you to read over the human subjects protection form with me. This is a document that explains any sort of risks you may face in participating in this study – there are none – and explains that you will receive money in exchange for your participation. (*Read over human subjects permission form*.)

If you agree to this and would like to participate, go ahead and sign and date at the bottom. If you would like an additional copy for your records, please let me know.

[Experimenter collects permission slips]

The first thing we need to determine is if any of you know any of the other participants in this room right now.

[If subjects know each other, keep track and make sure they are in the same group. If it is not possible to do this for whatever reason, give the subjects that make it impossible $5 and thank them for coming.]

Next, we are going to divide you into two groups: Group A, and Group B. Those of you sitting on the left hand side of the room will be in group A, those on the right hand side will be group B.

I will now assign an ID# to each member of group A. This number will be between 1 and 5.

[Assign each member of group A an ID#]

Will each member of Group A please stand and state their ID number. Group B, please do your best to remember the face and ID number of each member of group A, you will be paired with one of them shortly.

[Let Group A introduce themselves]

In just a moment we will ask the members of group B to move to a separate room at which point each member of group B will be randomly assigned to be the partner of a member of group A. Group B, this means that you will know the identity of your partner. Group A, you will not know the identity of your partner, but your partner will know your identity.

Group A, it will be up to you to determine the payments you and your partner will receive. Group B, you will have no say in the payments you and your partner will receive. Your payment will be determined by the actions of your partner in Group A. This process will involve the dividing of various sums of tokens each of which is worth ten or twenty cents ($0.10 or $0.20). The amount passed from the Group A member to their group B member may be multiplied by

some amount in determining payments. Group A members will be asked to make many such decisions, but only one will be selected at random to be carried out for actual money.

A second experimenter is presently waiting outside this room. Group B, if you would please follow him into the next room, you may bring your things with you. Thank you again for participating. Group A, stay here with me, we will begin once Group B has left the room.

II. Instructions to Group A (T1) – *These instructions were handed out and read to members of group A before they began practicing on their computers.*

In the first 10 settings, you have two piles of tokens – Left and Right – and the opportunity to pass as many of those tokens from either or both piles to your partner as you want. Each token is worth $0.10.
You choose how much you would like to pass by using the slide bar. You can tell how many tokens you have available to pass from each pile by the "Pile description" written above each slide bar.

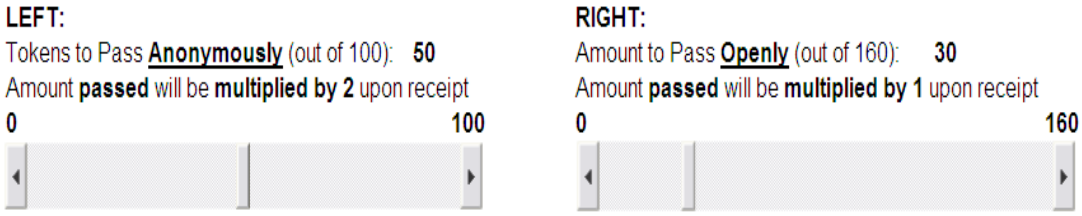

Why two piles? The Left and Right Piles differ in three ways:

1. They have different amounts of tokens available – here, the Left has 100, the Right has 160.
2. Tokens passed are multiplied by different amounts – tokens passed from the Left Pile in this case are multiplied by 2 when given to your partner (but not when received by you). Tokens passed from the Right Pile here are multiplied by 1. So, tokens passed from the Left pile are worth twice as much. Pay attention to these numbers. In some settings, tokens passed from one pile may be worth up to 10 times as much.
3. MOST IMPORTANTLY – Your partner does not know the Left Pile exists. They will receive tokens you pass from it, and they will be multiplied, but they will not know that they are coming from you. They will receive these tokens as a "thank you" payment. Your partner does know about the Right Pile, so you will receive recognition for tokens passed from the Right Pile.

Notice that in the above example you would be passing 50 tokens anonymously and 30 tokens openly. Since tokens passed anonymously are multiplied by 2 in this setting, your partner will receive 100 tokens ($10) as a "thank you" payment. Since tokens passed openly are multiplied by 1 in this setting. Your partner will receive 30 tokens ($3) openly from you.
This means that, in total, your partner will receive $13, but as far as they know, only $3 of this came from you. Also, it appears to them as though you are keeping 130 tokens (you had 160 to and passed 30), when in reality you're keeping 180 (130 openly and 50 anonymously). Keeping track of all these payoff calculations might be complicated, so we do it for you.

To help visualize your and your partner's payoffs from your perspective and their perspective (which is different from yours since they do not know the Left Pile exists) we provide you with two pie charts.

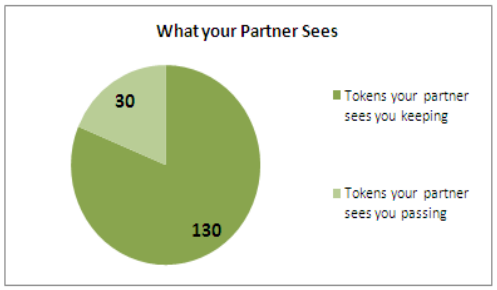

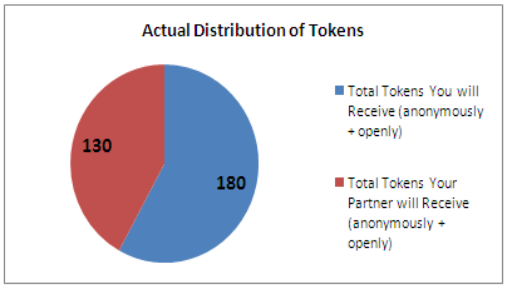

The pie chart on the Left shows what your partner sees. As explained above, they see you passing 30 tokens and keeping 130 tokens.
The pie chart on the Right shows what is actually happening. You are passing 130 tokens and keeping 180 tokens.

Each of the ten settings is different in that the number of tokens in each pile is different as is the amount by which tokens passed are multiplied. The one thing that does not change is that your partner will never be informed about the existence of the Left Pile.

Summary:

- You have two piles of tokens (Left and Right)
- Your partner does not know the Left Pile exists (you only receive recognition for actions concerning the Right Pile)
- Tokens are multiplied when passed (not kept) so it is important to remember that tokens passed from one pile may be more valuable to your partner than tokens passed from another pile.

Calculating your Payment:
ONE setting will be chosen at random (by rolling a 20 sided die) to be carried out for real money. Each setting has an equal chance of being chosen, so it is important to treat each setting as though it will be the one that is chosen.

To streamline the manner in which we calculate your payments, a "decision table" is provided at the bottom of each setting that looks like this.

If you are pleased with your decisions in this setting, enter the numbers below into the corresponding decision slip for Setting #EXAMPLE

ID#	Setting	A	B	C	D	E	F	G	H	I	J
0	Example	50	50	50	100	130	30	130	30	180	130

You will be given decision slips which look exactly the same (except without the numbers filled in), and your job is to fill in the numbers indicated in your decision table. Make sure you fill out the numbers for the appropriate setting or we won't be able to calculate your payment.

Combining the two piles, the two pie charts, and the decision tables, each setting looks like this:

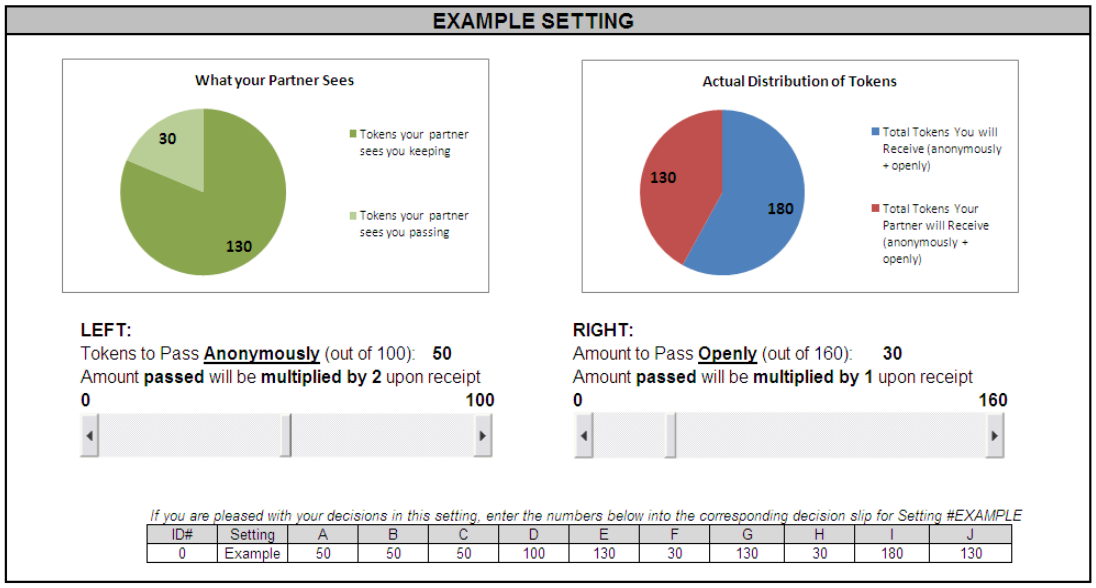

ID#	Setting	A	B	C	D	E	F	G	H	I	J
0	Example	50	50	50	100	130	30	130	30	180	130

Move over to your computers and we will practice using this together.

III. Instructions to Group B – *These were read to Group B by the second experimenter after group B had been moved into a separate room.*

Thank you again for participating in this experiment. As members of Group B you have two jobs to do. One is to take a short survey which will be explained shortly. The second is to sit back and wait while your partner in Group A determines how much money to share with you. I'll explain this first.

The first step in this process is matching you with a partner. If you recall, the people in group A were each assigned an ID# 1 through 5. I will now pass around a stack of cards, each with a number 1 through 5 written on it. Please pick one card to determine who you are going to be paired with.

[Experimenter circulates cards]

Your partner will never know which one of you they are paired with, though they do know that they have been paired with one of you and that *you* know the identity of your partner (and also that you know that they don't know your identity).

To determine your and your partners' payoffs, your partners are being asked to divide various piles of tokens between yourself and themselves in various settings. Each token is worth ten or twenty cents ($0.10 or $0.20). In these various settings, the total number of tokens available to be divided vary, as does the amount by which each token passed to you will be multiplied when we determine your payment. Your partner is aware of this multiplier as well as the total number of tokens and also of the fact that you are being informed about these things. Your partner is making many such divisions, but only one will be chosen at random from all of their decisions to determine your and your partners' payments. You will only be informed of the decision made by your partner which was selected – at random – to determine your payments.

There are two specific scenarios in which your partner will be dividing tokens and determining payments. These processes are best explained by example.

In the first group of scenarios (there are 10 of these total), your partner has the opportunity to divide various piles of tokens. Suppose that, in one such setting, your partner has the opportunity to divide a pile of 120 tokens and is made aware that any tokens passed to you will be multiplied by 1.5 when determining your payment (and tokens they choose to keep will be multiplied by 1). If your partner elects to pass 40 tokens to you and keep 80 for themselves, you will receive 60 tokens (40*1.5 = 60) or $6, and they will receive 80 tokens (120 – 40 = 80) or $8.

CHAPTER 3

CHOOSING BETWEEN MIXED AND PURE STRATEGIES IN A SIMPLE EXPERIMENT

Bryan Tomlin

University of California, San Diego

MAY 2010

Abstract

The results of an experiment studying the choices of subjects playing two simple variants of a 2x2, constant sum, strictly competitive, matching pennies game are presented. The experiment allows subjects to choose a mixed strategy in the literal sense: by selecting their preferred probability distribution over their action space. In a "best-of-three" treatment, the mixed strategies chosen by subjects are used to determine behavior in three rounds against the same opponent (with the player winning the majority of games winning some constant payoff). Average and modal subject behavior in both treatments conforms to that predicted by the Nash Equilibrium, though subjects in the baseline treatment are more than three times as likely to play a pure-strategy as are subjects in the best-of-three treatment. A preference model is proposed as a potential explanation for this behavior.

I. Introduction

Will experimental subjects elect to play mixed strategies when there is no possibility of repeated interaction and all strategies, including pure-strategies, are rationalizable? Can the propensity to play mixed strategies be manipulated by simply asking subjects to consider repeated interactions? To answer these questions we examine two experimental treatments of the 2x2, constant sum, strictly competitive, matching pennies game shown below:

The Matching Pennies Game[37]

		Column Player	
		Heads	**Tails**
Row Player	**Heads**	(1, 0)	(0, 1)
	Tails	(0, 1)	(1, 0)

To elicit subjects' mixed strategies in a non-repeated framework, we allow subjects to select their preferred probability distribution over their action space.[38] Subject select the probability (p) with which they would like to select Heads (H), and subjects who choose p=0 or p=1 are said to be playing a pure strategy. Those who choose otherwise are said to be playing a mixed strategy.

In a "Single Round" (SR) treatment, each row (column) player is randomly paired with a single column (row) player and asked to make a single selection of p. This value is then used to determine their choice of H or Tails (T) in a single game. Only one game is played. In a "Best of Three" (Bo3) treatment, each row (column) player is randomly paired with a single column (row) player and asked to make a single selection of p. This value is then used to determine their choices of H and/or T in three games against the same opponent. The winner of the majority of

[37] This perspective (i.e. the 2x2 simultaneous game table and the use of the words "heads" and "tails") is not the same as that presented to subjects (though they were asked to play this very game). Appendix A presents the game as viewed by subjects in this experiment.

[38] This experimental device was first utilized by Shachat (2002), in which he studied behavior in a 4x4 zero-sum stage-game repeated over 65 rounds.

these three games is the Winner.[39] The Winners do not earn more for winning 3 games rather than 2 games (Winners receive the same earnings regardless of "how" they win). To be clear, in the Bo3 treatment, subjects still only select one value for p which is then used to determine their action in three identical games of the form described above (so it is possible for subjects to play H in one game and T in another so long as $0 < p < 1$).

The SR and Bo3 treatments have the following in common:

- The payoff received for being the Winner is equivalent in both experimental treatments
- There is no potential for learning in either game since only one choice is made
- There is no potential for manipulating opponent beliefs since opponents never meet and choices are not revealed until they have been made (and only one choice is made)
- They both have a unique Nash Equilibrium of $p = 0.5$.
- *All* strategies (including p=0 and p=1) are rationalizable in both treatments.
- For any subject, if a strategy is optimal (for any reason) in the SR treatment, it should also be optimal in the Bo3 treatment. If they believe their strategy gives them the best chance of winning a single game, they should also believe it gives them the best chance of winning the best of three games.

Given this lack of learning and/or strategic manipulation, as well as the identical (unique) Nash Equilibrium and set of rationalizable strategies, we would expect subject behavior in SR and Bo3 to be very similar. Instead, we observe a significant shift in behavior away from pure strategies and towards the Nash Equilibrium: specifically, subjects in the SR treatment are more than three times as likely to play a pure strategy, while subjects in the Bo3 treatment are 60% more likely to play the Nash Equilibrium prediction of p=0.5.

[39] For clarity, the winner of a single game in the Bo3 treatment is referred to as the "winner" with a lower-case "w" while the winner of the majority of the three single games is referred to as the "Winner" with a capital "W."

Risk preferences over monetary gains alone are unable to explain these differences, though a theory of beliefs and preferences over the number of games won in the Bo3 treatment is proposed which is able to explain these results.

Section II discusses relevant literature, section III describes the treatments and the results, section IV proposes possible explanations for the observed behavior, and section V concludes.

II. Previous Literature

To study learning and how subjects adapt their strategies to information, Suppes and Atkinson (1960), and Malcolm and Lieberman (1965), generalize the mixed strategy being played in a stage game of a repeated MP experiment using the frequency of play of H or T both between subjects (within "rounds") or within subjects (across "rounds"). Since subjects in these experiments are confined to selecting one action per round (as opposed to a probability distribution over their actions) the assumption in these experiments is that the mixing observed in the data (i.e. the switching between actions across rounds) is the result of subjects determining the outcome of their own mental randomization.[40]

O'Neill (1987) ran a similar experiment (two-player, zero-sum) that incorporated a game with a unique mixed-strategy minimax solution and used aggregate data to find that aggregate subject behavior conformed rather closely to the minimax solution. Brown and Rosenthal (1990) conducted a review of O'Neill's data and found that data on the *individual* subject level was not consistent with minimax play. They found serial correlation in O'Neill's subject's choices,

[40]Ochs 1995 and Bloomfield 1994 allow subjects to select the number of times (N) over the next K rounds that they would like to play an element of their action space. The subjects are told that their computer will then randomly order these actions for the upcoming K rounds of play. While this approximates a mixed strategy in terms of the expected outcome in any one of the K rounds it does not provide a true mixed strategy. Shachat 2002 discusses in detail why neither of these experiments should be thought of as eliciting subjects' actual mixed strategies. For this discussion, we refer you to his work.

correlation which Shachat (2002) (henceforth "Shachat") explains may be the result of subjects playing repeated game strategies that deviated from the minimax solution, or the result of subjects being cognitively unable to a construct a random sequence of i.i.d. actions. Shachat avoids these problems by conducting an experiment incorporating a "mixed strategy device," which allows subjects to directly "plug in" their preferred mixed strategy (this is a literal mixed strategy in that the subjects select a probability distribution over their action set). This bypasses the cognitively difficult task of mentally generating a sequence of independent, random, actions, such that any serial correlation which may exist in the observed sequence of actions is entirely due to a deviation from minimax play. Shachat finds that subjects do, indeed, deviate substantially from minimax play.

The above referenced experiments all utilize repeated stage games to study mixed strategy play.[41] However, the basic theoretical notion of a mixed strategy does not require repeated interaction or even a multi-staged game. The experiment presented here borrows Shachat's "mixed strategy device" to answer the simple question: do subjects elect to play a mixed strategy in a non-repeated game? Furthermore, since mixed strategies are often motivated by repeated interactions (i.e. the above mentioned literature), perhaps the very notion of repeated interactions and mixed strategies are somehow linked. To study this, the Bo3 treatment introduces a weak repeated interaction[42] that does not introduce risk preferences or added strategic consideration. As such, any observed difference in play between these two treatments should be the result of this pseudo-difference in repeated interaction.

[41] Not for the sake of simply studying mixed strategies, but instead to answer questions about learning, strategic adaptation, solution concepts etc. They do not attempt to answer the simple question being motivated here.

[42] "Weak" in the sense that the stage game is not being repeated at all (since only one choice is made).

III. Experimental Description and Results

Subject Recruitment and Payment

Subjects were University of California, San Diego students enrolled in three different Econ 1 courses. Econ 1 (introductory economics) is a course which must be successfully completed in order to enroll in UCSD's Game Theory course. Elementary game theoretic concepts are occasionally introduced in this course, though the students surveyed for this research had not yet been introduced to these concepts.

Subjects were assigned to treatment (Single Round (SR) or Best of Three (Bo3)) depending on which lecture they were enrolled in. Subjects were randomly assigned to the role of either the Row or Column Player. All instructions were read aloud as shown on the full instructions which can be found in Appendix A.

To avoid steering responses and to better ensure understanding, subjects were asked to provide the percent chance with which they would like to play *both* of their strategies. Subjects whose responses did not sum to 100% were dropped. Of the 191 subjects surveyed, only two provided responses which did not sum to 100 (both the in the Bo3 treatment): one subject's responses summed to 80% (40%+40%), the other's to 130% (80% + 50%). This left 189 useable subjects (96 in SR, 93 in Bo3).

Treatment Description

In both treatments subjects are paired anonymously and never meet their opponent. Winning players received entry into a drawing for a $20 prize. One prize was given out for each treatment. The expected value of this prize was therefore $0.42 in SR and $0.43 in Bo3. The descriptions below are replicated from the introduction for ease of reference.

<u>Single Round Treatment (SR)</u>

In SR each row (column) player is randomly paired with a single column (row) player and asked to make a single selection of p (the probability with which they would like to play Heads). This value is then used to determine their choice of Heads (H) or Tails (T) in a single game. Only one game is played. The winner receives entry into the drawing for $20 described above.

Best of Three Treatment (Bo3)

In Bo3 each row (column) player is randomly paired with a single column (row) player and asked to make a single selection of p. This value is then used to determine their choices of H and/or T in three games against the same opponent. The winner of the majority of these three games is the Winner. The Winners receives entry into the drawing for $20 described above. To be clear, in the Bo3 treatment, subjects still select only one value for p which is then used to determine their action in three identical games of the form described above (so it is possible for subjects to play H in one game and T in another so long as $0 < p < 1$).

Results

Table 1 describes subject behavior in SR and Bo3. Figure 1 provides the distribution of p for both treatments.

[Table 1 – Here]

[Figure 1 – Here]

A t-test of the difference in the mean value of p between treatments finds average probability weights on heads to be statistically indistinguishable between experimental conditions ($p=0.536$). However, a test of differences in the percentage of subjects choosing to play a pure-strategy in SR (16.7%) relative to the percentage of subjects choosing to play a pure-strategy in Bo3 (5.4%) shows that subjects in SR are significantly more likely to play a pure-

strategy (p=0.014). Furthermore, subjects in Bo3 are significantly more likely to play the Nash Equilibrium (probability of heads = 0.5) strategy than subjects in SR (26.9% vs. 16.7%, p=0.089).

IV. Possible Explanations

Explanation 1 – Considering Repeated Interaction

One possible explanation for these results is that subjects in the Bo3 setting are being primed to think about repeated interactions. By being placed in a strategic framework where their payoff depends on the result of seemingly repeated interactions (players are not actually interacting repeatedly since they only select a single action), subjects may be more likely to consider how they would behave were there to be repeated interactions. As observed by Shachat (2002), subjects provided with a mixed strategy device elect to play different mixed strategies in different rounds of a repeated stage game, even when their opponent does not observe their chosen probability distribution. The observed increase in the number of p=0.5 choices and decrease in the number of p=0 and p=1 choices in the Bo3 treatment relative to the SR treatment may represent subjects "averaging" their choices over three repeated games. For example, suppose the stage game were to be repeated three times and a subject would have chosen p=0 in the first stage, p=1 in the second stage, and p=0.5 in the third. This subject, if placed in the SR treatment might prefer to simply play any one of these three strategies. However, if placed in the Bo3 treatment, they might prefer to play a mixture of all three of their desired strategies at once, averaging over all three, and selecting p=0.5. In other words, subjects may be treating the Bo3 treatment as three repeated stage games by selecting a compound mixed strategy composed of a mixture of whatever strategies they would have chosen over three repeated stage games.

Another possible behavioral explanation involves the "Law of Small Numbers." Not wanting to be predictable, subjects may be attempting to mentally simulate the amount of temporal mixing which will take place in the three outcomes of the Bo3 treatment and believe that a choice of p=0.5 will produce that mix better than p=0 or p=1. People are inclined to overestimate the amount of switching between outcomes which will occur in a binomial sequence, and believe that "short" segments of a random sequence should reflect the true distribution of outcomes. If subjects are attempting to be as random as possible in the Bo3 treatment, they may therefore consider a sequence with two H and one T or one H and two T to be "more random" than a sequence of all H or all T. Since p=0 and p=1 surely result in all H or all T, they will be considered "less random" and therefore more predictable than p=0.5.

Explanation 2 – A "Shutout" Penalty and "Perfect" Bonus

To explain the increase (decrease) in the number of p=0.5 (p=0 or 1) choices in the Bo3 treatment relative to the SR treatment, suppose that some players are trying to decide between playing p=0, p=1, or p=0.5. Furthermore, suppose that players believe that their opponents are either playing p=0.5, p=0 or p=1, and that the proportion of opponents playing p=0.5 is $(1 - \alpha)$, the proportion playing p=0 is $\alpha/2$, and the proportion playing p=1 is $\alpha/2$ (so α of opponents are believed to be playing a pure strategy and the remaining $1 - \alpha$ are believed to be playing the Nash Equilibrium strategy). Any choice of p is rationalizable given these beliefs.

Let $U_{1,SR}$ and $U_{0,SR}$ represent the utility of winning and losing respectively in the SR treatment. Let U_3, U_2, U_1, and U_0 represent the utility of winning 3, 2, 1, and 0 games respectively in the Bo3 treatment.

The expected utility of playing a pure strategy in the SR treatment is given by:

(1) $(\alpha/2)[U_{1,SR} + U_{0,SR}] + (1 - \alpha)[U_{1,SR}/2 + U_{0,SR}/2] = (U_{1,SR} + U_{0,SR})/2$

The expected utility of playing p=0.5 in the SR treatment is given by:

(2) $(\alpha/2)[U_{1,SR}/2 + U_{0,SR}/2] + (\alpha/2)[U_{1,SR}/2 + U_{0,SR}/2] + (1 - \alpha)[U_{1,SR}/2 + U_{0,SR}/2] = (U_{1,SR} + U_{0,SR})/2$

As (1) is equivalent to (2), players will be indifferent between playing a pure strategy and p=0.5 in the SR treatment, regardless of the values of $U_{1,SR}$ and $U_{0,SR}$, and regardless of the player's risk preferences. This result is supported by the data which finds the exact same number of subjects (16.7%) to be playing p=0.5 as p=0 or p=1 in the SR treatment.

The expected utility of playing a pure strategy in the Bo3 treatment is given by:

(3) $(\alpha/2)U_3 + (\alpha/2)U_0 + (1 - \alpha)[0.125U_3 + 0.125U_0 + 0.375U_1 + 0.375U_2]$

The expected utility of playing p=0.5 in the Bo3 treatment is given by:

(4) $\alpha[0.125U_3 + 0.125U_0 + 0.375U_1 + 0.375U_2] + (1 - \alpha)[0.125U_3 + 0.125U_0 + 0.375U_1 + 0.375U_2]$

$= [0.125U_3 + 0.125U_0 + 0.375U_1 + 0.375U_2]$

Players will choose to play either p=0.5 or a pure strategy depending on the relative values of conditions (3) and (4). Combining (3) and (4) and simplifying, we see that:

(5) If $U_1 - U_0 > U_3 - U_2$, the subject will prefer p=0.5 to a pure strategy

(6) If $U_3 - U_2 > U_1 - U_0$, the subject will prefer a pure strategy to p=0.5

(7) If $U_3 - U_2 = U_1 - U_0$, the subject will be indifferent between p=0.5 and a pure strategy.

If a player's utility is entirely an increasing function of their earnings, then $U_3 = U_2 > U_1 = U_0$, since they will only receive payment for winning two or three games. In this case, $U_3 - U_2 = U_1 - U_0$, implying that the subjects will be indifferent between p=0.5 and a pure strategy. However, the portion of players choosing pure strategies in the Bo3 treatment is significantly smaller than the number of players choosing pure strategies in the SR treatment (p=0.014). It is therefore unlikely that players are in fact indifferent between playing p=0.5 and playing a pure strategy in the Bo3 treatment.

Suppose that players receive utility not only from their earnings, but also from winning more games, such that $U_3 > U_2 > U_1 > U_0$. In this case, we can rethink the values of $U_3 - U_2$, $U_2 - U_1$, and $U_1 - U_0$ as follows:

Utility Difference	**Interpretation**
$U_3 - U_2$	"Perfect Bonus" – The added utility associated with winning all three games. Unassociated with a change in wealth.
$U_1 - U_0$	"Shutout Penalty" – The added utility associated with not being shutout (not being winless). Unassociated with a change in wealth.
$U_2 - U_1$	"Winner Bonus" – The added utility associated with winning the majority (rather than the minority) of games *and* receiving payment.

Given this interpretation, players will choose to play a pure strategy in Bo3 if their Perfect Bonus is greater than their Shutout Penalty, and will choose to play p=0.5 if their Shutout Penalty is greater than their Perfect Bonus, *regardless of the size of their Winner Bonus*. Figure 2 provides graphical illustration of utility functions which could represent player preferences. Notice that it is not necessary to assume strictly concave or convex utility of winning preferences (as illustrated by the center column in Figure 2) to arrive at this result.

[Figure 2 – Here]

As shown in section III, subjects play p=0.5 more frequently than pure strategies in the Bo3 treatment. This can be interpreted as subject preferences being best described by condition (5) above (the bottom row of Figure 2). There are significantly more subjects who try to avoid being shut out then there are subjects attempting to shut-out their opponent.

V. Conclusion

The proposed explanation for the result of this simple experiment – that subjects dislike being "shutout" more than they enjoy shutting out their opponent – is intuitively pleasing. The existence of "mercy rules" in little league sports, or fans' disfavor for teams who "run-up the score" seem to validate this notion. This explanation also coexists nicely with the fact that some teams consider it to be worth risking injury to obtain a "perfect season" while others do not.[43]

The explanation that subjects may play differently when they are asked to consider repeated interactions is also intuitively pleasing. The mixing device utilized by this experiment offers subjects the opportunity to choose not to "do the same thing" in all three games of the Bo3 treatment. Subjects may choose to vary their behavior across this small number of games to avoid being predictable in accordance with the Law of Small Numbers.

Whether either of the proposed explanations are in fact, *the* explanations for the above described behavior cannot be determined from this experiment. Additional explorations (and explanations) of this type of behavior would be very interesting.

[43] The pursuit of a "Perfect Season" in a sport such a professional football is an interesting example of this phenomenon. Teams who have already won enough games to secure home-field advantage throughout the playoffs face the decision of resting their most valuable players in preparation for the post-season or risking injury to these players in "pursuit of perfection" (i.e. an undefeated record). Unlike the experiment presented in this paper, there are clear costs associated with the decision to pursue a perfect season in the NFL which make the decision to do so all the more uncommon.

CHAPTER 3 – BIBLIOGRAPHY

Bloomfield, R. “Learning a mixed strategy equilibrium in the laboratory,” *Journal o. Economic Behavior Organization*, 1994

Brown, J., and Rosenthal, R. “Testing the Minimax Hypothesis: A re-examination of O’Neill’s game experiment,” *Econometrica*, 1990

Camerer, C., and Ho, T. "Experience Weighted Attraction Learning in Normal-Form Games." *Econometrica*, 1998

Cheung, Y. and Friedman, D. "Individual Learning in Normal Form Games: Some Laboratory Results." Mimeo, University of California, Santa Cruz, 1995

Chiappori, P., Levitt, S., and Groseclose, T., “Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer,” *The American Economic Review*, 2002

Cox, J., Shacht, ., and Walker, M. “An Experiment to Evaluate Bayesian Learning of Nash Equilibrium.9"
Mimeo, University of Arizona, 1995

Crawford, V. "'An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Co-ordination." *Games and Economic Behavior*, February 1991

Crawford, V. " Adaptive Dynamics in Coordination Games." Mimeo, University of California, San Diego, 1992.

Erev, I., and Roth, A., “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *The American Economic Review*, 1998

McCabe, K., Mukherji, A., Runkle, D., “An Experimental Study of Information and Mixed-Strategy Play in the Three-Person Matching-Pennies Game,” *Economic Theory*, 2000

Ochs, J., “Games with Unique, Mixed Strategy Equilibria: An Empirical Study,” *Games and Economic Behavior,* 1995

O’Neill, B. “Nonmetric Test of the Minimax Theory of Two-Person Zerosum Games,” *Proceedings of the National Academy of Sciences USA*, 1987

Pearce, D., “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 1984

Rabin, M. “Inference by Believers in the Law of Small Numbers,” *Quarterly Journal of Economics*, 2002

CHAPTER 3 – TABLES AND FIGURES

Table 1: Subject Behavior in the Single Round and Best-of-Three Matching Pennies Game Treatments

	Single Round (SR)	Best of Three (Bo3)
Number of Observations	96	93
Mean p(H)	0.495	0.518
Median p(H)	0.50	0.50
Mode p(H)	0.50	0.50
Number of p(H) Є {0, 1}	16 (16.7%)	5 (5.4%)
Number of p(H) = 0.50	16 (16.7%)	25 (26.9%)

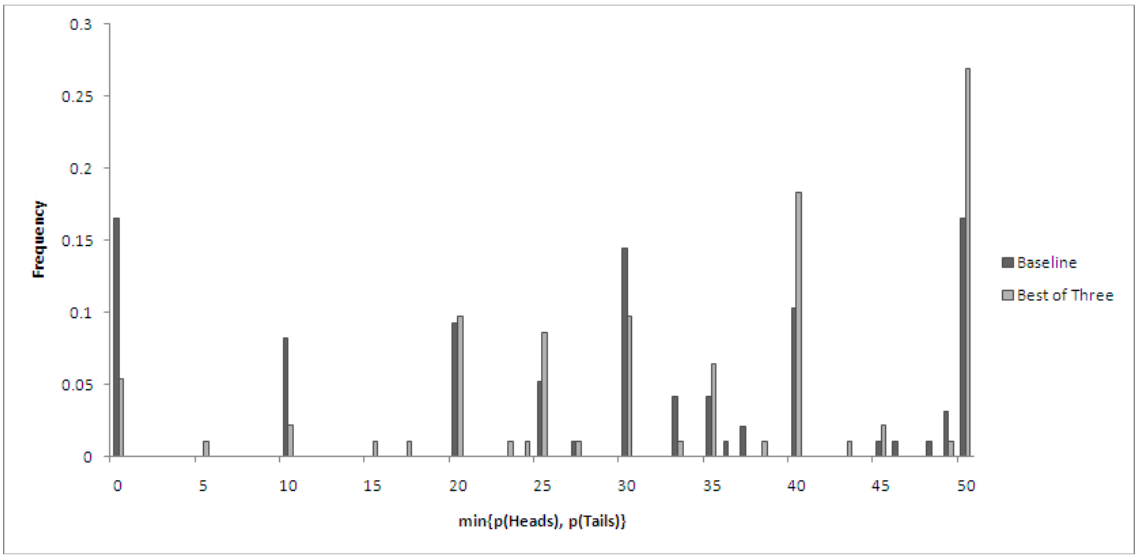

Figure 1: Distribution of the Frequency of min{p(Heads), p(Tails)} in the Single Round and Best-of-Three Treatments

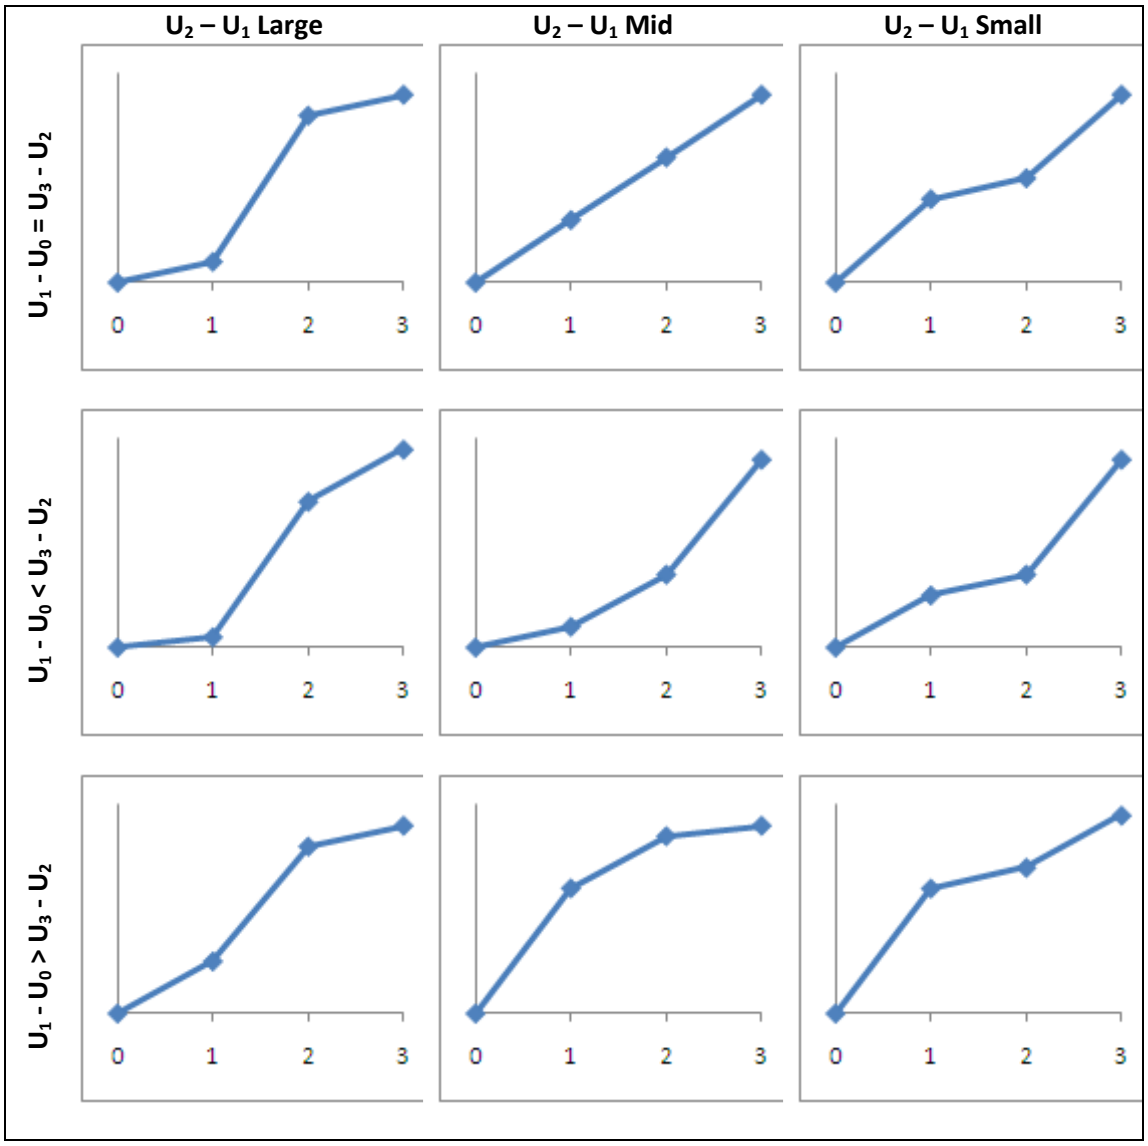

Figure 2: Possible Utility Values (Vertical Axis) as a Function of Number of Games Won (Horizontal Axis) in the Best of Three (Bo3) Matching Penny Game Treatment

CHAPTER 3 – APPENDIX A – INSTRUCTIONS TO SUBJECTS

The instructions administered to subjects are provided below. Words written inside of brackets were not provided on the instructions and are placed there for the sake of the reader of this paper. Only instructions to player 1 are provided.

[Single Round Treatment:]

To identify yourself in case you win, write the last 5 digits of your PID here: _____

You are **PLAYER 1**

- You will be playing against Player 2, a randomly selected individual playing the same game who you will never meet.
- Player 1 chooses either UP or DOWN
- Player 2 chooses either LEFT or RIGHT
- This game will only be played once.
- The winner, as determined by the table below, will be entered into a drawing to win $20. You will only be notified that you have won if you are selected as the winner of the $20.

Player 1 Chooses	Player 2 Chooses	Winner
UP	LEFT	Player 1
UP	RIGHT	Player 2
DOWN	LEFT	Player 2
DOWN	RIGHT	Player 1

Please specify the *percentage chance* with which you would like to play "UP" and "DOWN." We will use a random number generator to make your decision according to the percentages you specify:

I would like to play UP with a _____% chance.
I would like to play DOWN with a _____% chance.
Each number must be no less than 0 and no greater than 100. These numbers must sum to 100.

Examples:

- If you play UP with a 20% chance and DOWN with an 80% chance, you will be four times more likely to play DOWN than UP.

- An option played with a 100% chance is played for certain (and an option played with a 0% chance is never played).

[Best of Three Treatment:]

To identify yourself in case you win, write the last 5 digits of your PID here: _____

You are **PLAYER 1**

- You will be playing against Player 2, a randomly selected individual playing the same game who you will never meet.
- Player 1 chooses either UP or DOWN
- Player 2 chooses either LEFT or RIGHT
- You will play this game against the *same opponent* three times.
- Whoever wins more games (as determined by the table below) will be entered into a drawing to win $20. You will only be notified that you have won if you are selected to win the $20.

Player 1 Chooses	Player 2 Chooses	Winner
UP	LEFT	Player 1
UP	RIGHT	Player 2
DOWN	LEFT	Player 2
DOWN	RIGHT	Player 1

Please specify the *percentage chance* with which you would like to play "UP" and "DOWN." We will use a random number generator to make your three decisions (one for each game) according to the percentages you specify (so it is possible to have UP chosen in one game and DOWN in another so long as neither % chance is set to 0)

I would like to play UP with a _____% chance.
I would like to play DOWN with a _____% chance.

Each number must be no less than 0 and no greater than 100. These numbers must sum to 100.

Examples:

- If you play UP with a 20% chance and DOWN with an 80% chance, you will be four times more likely to play DOWN than UP.

- An option played with a 100% chance will be played for certain in all three games (and an option played with a 0% chance is never played).