

Pathways of Persuasion*

Lucas Coffman[†]
Ohio State University

Paul Niehaus[‡]
UC San Diego

September 1, 2014

Abstract

While economic theories of persuasion emphasize appeals to self-interest, others posit an important role for other-regard. For example, a salesperson might describe product features but also try to build rapport. We study these two mechanisms within a simple experimental framework in which buyers talk with sellers incentivized to raise their valuations for objects. We find that sellers benefit from communication despite their material conflict of interest. They affect the distributions of both buyers' self-interest and other-regard, though only the former effect is mean positive. A minority of sellers do strategically target other-regard, however, and outperform their peers. The mechanics of persuasion vary with the nature of the good being sold and with buyer-seller homophily – particularly gender match. Surprisingly, variation in persuasion is driven less by who is selling than by who is buying.

*We thank seminar participants at UC Berkeley Haas School, Brown University, SITE, UC San Diego Rady School, UC San Diego, the Florence Workshop and ESA Santa Cruz for helpful comments. An earlier version of the paper circulated under the title “Interpersonal Influence.”

[†]Department of Economics, Ohio State University, 1945 North High Street #410, Columbus, OH 43210. coffman.155@osu.edu.

[‡]NBER and Department of Economics, University of California at San Diego, 9500 Gillman Drive #0508, San Diego, CA 92093-0508. pniehaus@ucsd.edu.

1 Introduction

People often use communication to *persuade* others to do what they want: politicians and voters, job candidates and employers, salespeople and customers, fundraisers and donors, CEOs and investors, and so on. One interesting feature of these interactions is that the “sender” might persuade the “receiver” in different ways. The approach emphasized within economics is to convince the receiver that a given action is in her own *self-interest*.¹ For example, a salesperson might describe appealing features of a product. Writers outside of economics, on the other hand, often emphasize the importance of being likeable, building rapport, and so on. The cover of an early edition of Dale Carnegie’s self-help classic “How to Win Friends and Influence People” promised, for example, to reveal “the six ways of making people like you” (Carnegie, 1936). Evidence from social psychology suggests that “people prefer to say yes to individuals they know and like” (Cialdini, 1993). In short, persuasion could also work through manipulation of *other-regard*.

In this paper we bring these two mechanisms – self-interest and other-regard – together to quantify their roles in persuasion. To do this we create experimental situations in which there is scope for one party to persuade another. For concreteness we focus on a buyer-seller paradigm: a buyer must decide what he is willing to pay for an object, and it is common knowledge that the seller will receive a commission if he buys, creating a potential conflict of interest.² We then allow for free-form conversation between the two and measure the impact this has on the buyer’s valuation. We intentionally do *not* constrain or vary the message space available to subjects in order to allow sellers scope to use any strategy they like.

We measure the effects of this communication and decompose them into changes in other-regard and (perceived) self-interest. Specifically, we elicit buyers’ willingness to pay for objects before and after conversations and also benchmark these changes against the effects of a placebo treatment in which buyers simply had time to think. We then decompose willingness-to-pay into two components: the amount they are willing to pay for the object when the seller does *not* receive a commission, and the *additional* (possibly negative) amount they are willing to pay when she does.³ Changes in the first component capture effects on self-interest, while changes in the second capture effects on other-regard.⁴

We complement our revealed preference measures with two additional data sources. First, we collect self-reported data from buyers and sellers before and after conversations on their

¹See Crawford and Sobel (1982), Olszewski (2004), Mullainathan et al. (2008), Kamenica and Gentzkow (2011), Dziuda (2011), and Che et al. (2013), among many others, for theoretical results on when and how such persuasion is possible.

²See Loewenstein et al. (2011) for evidence on the effects of disclosing such conflicts.

³This difference can be thought of as a measure of altruism. See Yamamori et al. (2008), Andreoni and Rao (2011), and Greiner et al. (2012) for evidence on the effects of communication in a pure dictator game. See Charness and Gneezy (2008) for effects of removing anonymity, Landry et al. (2006) on face-to-face interaction, and Leider et al. (2009) on social distance.

⁴Our distinction between self-interest and other-regard is thus related to but distinct from the (helpful) distinction that DelaVigna and Gentzkow (2009) draw between belief-based and preference-based persuasion. To highlight the difference, beliefs could affect either perceived self-interest (“it seems like a great product”) or other-regard (“he seems like a great guy”). Our distinction is also different from those in dual-process theories from psychology (e.g. Petty and Cacioppo (1986)) which contrast different modes of information processing.

plans for and interpretation of the interaction. Second, we transcribe and code all conversations. To discipline this content analysis we use a novel competitive mechanism: we asked each of seven leading researchers on persuasion across economics and psychology to propose the five metrics they thought would best predict persuasion, and promised to reveal the winner.⁵ This arguably gives us a snapshot of the state-of-the-art in forecasting while obviating the usual concerns about data mining. To address the concern that coding is inevitably subjective we draw our coders from the same subject pool as the original buyers in our experiment, ensuring that their subjective interpretations are statistically identical to those of actual participants.

Our experimental design has the advantage of adaptability to many contexts – one could vary the identities of participants, nature of the decision problem, communication channel, and so on. We necessarily focus on one particular context as a starting point, but also introduce some limited experimental variation to suggest hypotheses for future work. Within each experimental session we repeat the procedure above four times, rematching buyers and sellers to discuss four different objects for sale. This structure gives us panel data with which to study individual effects, and also lets us examine how variation in characteristics of the object being sold affects persuasion. We chose objects to create variation in the degree of information asymmetry between buyers and sellers. Two are consumer electronics devices with which sellers had (on average) more familiarity than buyers. The other two are simpler and fully specified: a lottery ticket with known payoff distribution and a post-dated check. We view the latter two as especially challenging objects to sell since a buyer who understands her own time and risk preferences should have immutable demand for them.

We find that seller communication does affect buyer behavior in spite of their misaligned material interests. Buyers change their valuations by an (absolute) average of \$49 following conversations, significantly higher than the \$24 average absolute change in placebo treatments. These valuation changes include individuals who would not have changed their valuations at all under placebo: 56% of valuations change after conversations compared to 31% in placebo. Moreover, sellers benefit on average from communication. The mean change in buyers' valuations with commissions is \$14 and significantly different from 0, compared to an insignificant mean change of \$3 in the placebo condition. Buyers thus do not interpret sellers' statements as cheap talk, despite knowing about their conflict of interest.

How do sellers persuade? Changes in self-interest and in other-regard both account for sizeable shares of overall variation in persuasion. Formally, conditioning on other-regard reduces the overall variation of persuasion by 16%, compared to 29% for (perceived) self-interest. Yet only self-interest changes on average: the mean change in buyers' valuations without commissions is \$20 and significantly different from 0, while the mean change in the incremental valuation of commissions is negative (−\$3) and insignificant. The “average” seller thus benefits from persuasion primarily because of buyer self-interest.

While these results suggest that other-regard does matter, they also raise the question whether sellers have any ability to *proactively* endear themselves, or whether buyers simply like some sellers more than expected and others less. To examine this we categorize conversations based on whether the seller's self-reported strategy prior to the conversation emphasized self-

⁵We congratulate Joel Sobel, whose measures predicted 18% of the variation in persuasion.

interest, other-regard, or was simply unclear. A minority (19%) of sellers do explicitly plan to manipulate other-regard, with strategies such as “get the customer on my side,” “build a rapport,” or “remember the person’s name! People love when others remember their names.” These sellers are significantly more likely than their peers to have tangential conversations that do not focus on the actual decision at hand, and also significantly more likely to use specific tactics thought to affect other-regard (for example, using the buyer’s first name).

Strikingly, sellers focused on other-regard out-perform their peers by a wide margin. They succeed in their stated objective, shifting other-regard by \$26 more than their peers on average. At the same time they do no worse than their peers at shifting buyers’ perceived self-interest. The upshot is that they benefit nearly three times as much from communication as other sellers. Even comparing conversations within-seller we find that the same seller does significantly better when targeting other-regard than when not. In short, manipulating other-regard appears to be a highly effective strategy for those who use it, though not widely used.

We also find evidence that both the extent and the mechanics of persuasion vary with the good being sold. Comparing consumer electronics goods to the abstract lottery and post-dated check, we find that sellers are \$15 more effective at persuading buyers to increase their valuations for the electronics. This difference aggregates even larger compositional differences: mean persuasion through self-interest is \$28 higher for the electronics, but persuasion through other-regard is \$13 lower. While stylized, these data suggest that which model of persuasion is “right” may vary predictably with the nature of the underlying decision problem, with self-interest playing a larger role where information is more asymmetric.

Finally, we examine the role of homophily (i.e. similarity between buyer and seller). We examine two specific traits in our data – gender and nationality – as well as an omnibus measure of buyers’ perceptions of similarity with the seller. We find strong evidence that homophily matters: a 1 s.d. increase in perceived similarity is associated with a \$18 increase in persuasion (a little higher than mean persuasion). Interestingly, this is driven by significant increases in both other-regard and self-interest, suggesting that homophily both directly affects liking and also makes communication more credible. This is particularly evident when we focus on gender: sellers raise valuations by \$22 more on average when matched with buyers of the same gender, and this effect is driven entirely by self-interest.

The results thus far, which describe influence for average participants, are most relevant for applications in which persuasion is a diffuse activity: the spread of network goods or a viral marketing campaign, for example. For other applications, however, it might be more important to understand what the *most* persuasive people do, as these individuals may be more likely to select into persuasion-intensive activities such as sales, negotiations, etc. We already know that sellers who target other-regard out-perform, which suggests that other-regard may play a larger role in the arsenal of expert persuaders. To let the data identify further patterns we exploit its panel structure, estimating the distribution of participant fixed effects. To our great surprise, variation in persuasion is driven less by sellers than by buyers. We cannot reject the null that all seller effects are identical, but do so resoundingly for buyers. In fact, buyer fixed effects explain more of the variation in our data than do seller fixed effects, product fixed effects, and order effects *combined*, and nearly as much as the content analysis recommended by our most successful expert predictor. From a forecasting perspective,

understanding what makes someone persuasive may thus be less important than understanding what makes someone *persuadable*.

Our results have broad implications for the literature on persuasion. The extensive theoretical canon on persuasion through self-interest would benefit from complementary work formalizing persuasion through other-regard, which remains under-theorized. Similarly, testable theories of persuadability could shed light on why some people are more receptive to persuasive messages than others and what consequences this has for the performance of markets. The results also beg a number of obvious empirical follow-up questions – for example, how and why do experts perform better than amateurs? What are the returns to face-to-face interaction (and do they justify business travel budgets)?

Our paper also contributes to the literature on the determinants of other-regard. A number of studies have documented the role of environmental factors such as proximity (Marmaros and Sacerdote, 2006), social distance (Leider et al., 2009), group interaction (Feigenberg et al., 2010; Shue, 2013), identification of partner (e.g. Bohnet and Frey (1999) and Small and Loewenstein (2003)), or the availability of communication (e.g. Andreoni and Rao (2011), Yamamori et al. (2008), and Greiner et al. (2012)). Our paper is part of a smaller literature showing how economic incentives affect the formation of social preferences *conditional* on such factors.⁶ Our results on homophily, meanwhile, contribute to the growing literature documenting its importance for a wide range of outcomes including trust (Glaeser et al., 2000), social learning (Golub and Jackson, 2012), and network formation (Currarini et al., 2009).

Finally, our results suggest a new perspective on bargaining. We view our design as a natural complement to standard bargaining designs: while these typically fix an objective value for the object being transacted and let participants negotiate price, we fix price (stochastically) and let participants discuss the subjective value of the object. Of course, many real-world situations combine elements of both designs. It would therefore be interesting to see how canonical results from the bargaining literature (e.g. the importance of fairness) interact with our results in a hybrid setting.

The rest of the paper is organized as follows. Section 2 describes the experimental design; Section 3 describes recruitment and participant characteristics; Section 4 presents the main results; and Section 5 summarizes and offers concluding thoughts.

2 Experimental Design

Our goal in designing the experiment was to create relatively natural and relevant settings for persuasion that also enabled us to capture rich data (including conversations themselves). We focus on a buyer-seller context because of its familiarity to subjects (as well as general economic importance). We organize the experiment as a series of interactions in each of which (1) a buyer submits an initial valuation for a good, (2) the buyer is randomly matched to a seller, (3) the buyer and seller have a free-form telephone conversation lasting at most 10 minutes, and (4) the buyer then submits an updated valuation. It is common knowledge that the seller’s (expected) payoff increases in the buyers’ willingness to pay, creating an incentive

⁶See for example Glaeser et al. (2002) who document evidence of social capital investment.

Table 1: Eliciting Willingness to Pay, Flat Wage

Please select which alternative you would prefer	
<input type="radio"/> I would prefer the new iPad	<input type="radio"/> I would prefer \$500
<input type="radio"/> I would prefer \$20 and the new iPad	<input type="radio"/> I would prefer \$500
⋮	⋮
<input type="radio"/> I would prefer \$480 and the new iPad	<input type="radio"/> I would prefer \$500

Each row is a separate question, and buyers answered all questions before proceeding with the experiment.

for her to persuade the buyer.

We define the difference between a buyers valuation before and after the phone call as persuasion. Identifying persuasion within-subject in this way maximizes our statistical power, which will facilitate analyzing specific determinants of persuasion. To address the concern that valuations might change over time even without conversation we also conduct a placebo arm, described below. Note that sellers were not given any information about buyers, including their ex ante or ex post valuations.

To decompose persuasion into effects working through self-interest and other-regard, we elicit buyers' valuations in two different ways: once for the object alone and once including a \$500 commission for the seller. To isolate self-interest, we elicit the buyers' valuations for the good alone. The buyer makes a series of incentivized decisions, choosing between \$500 or the object plus \$X. As the instructions explain to the subjects, the price of the good is thus \$500-X for that decision. We vary X from \$0 to \$480 in increments of \$20. (See Table 1.) We define the buyers valuation for the good as \$500 less the smallest X such that the buyer chooses the good plus \$X over \$500, and define changes in this quantity due to conversation as persuasion through the self-interest channel.

To measure the buyer's altruism towards the seller, we also elicit the buyer's valuation for the good with the seller receiving a \$500 commission. Immediately following the valuation elicitation for the product alone, the buyer makes a sequence of decisions between either \$500 for herself and \$0 for the seller, or the object and \$X for herself and \$500 for the seller. (See Table 2.) The only change from the first elicitation is that the seller now receives \$500 if a purchase is made; thus, how much more (or less) the buyer is willing to pay under commission captures how much she values the seller earning \$500.

We incentivize the valuation decisions in the following way: For every fifty buyers that participated in the experiment, a computer lottery held roughly one week later randomly chooses one to win \$500. The lottery winner also has one of her decisions across the entire experimental session randomly chosen to be implemented. For example, if a commission question is chosen where she chooses the good, she is given the good and \$X, and the seller is paid \$500. Since any question might be implemented, the buyer's incentive is to report truthfully her preferences to any one question. The seller's material incentive is to increase the buyer's WTP for the commission questions. The higher is the buyer's WTP, the more questions the buyer will select the alternative in which the seller receives a commission. Since

Table 2: Eliciting Willingness to Pay, Commission

Please select which alternative you would prefer	
<input type="radio"/> I would prefer the new iPad for me and \$500 for the seller.	<input type="radio"/> I would prefer \$500 for me and \$0 for the seller.
<input type="radio"/> I would prefer \$20 and the new iPad for me and \$500 for the seller.	<input type="radio"/> I would prefer \$500 for me and \$0 for the seller.
⋮	⋮
<input type="radio"/> I would prefer \$480 and the new iPad for me and \$500 for the seller.	<input type="radio"/> I would prefer \$500 for me and \$0 for the seller.

Each row is a separate question, and buyers answered all questions before proceeding with the experiment.

we randomly choose one question for payment, a higher WTP for the bonus questions increases the likelihood the seller is paid her bonus. This is explained to the seller in her instructions (All experimental instructions and questionnaires available on request).

Our basic design is highly flexible, with scope to accommodate different communication technologies, different participant types, and different objects for sale (or even alternative decision problems). While we certainly do not fully exploit this potential, we do introduce some variation in participant matching and in the set of objects for sale. Specifically, every subject participates in four buyer-seller interactions, always in the same role; every interaction is with a new counterpart and with a different object on sale.

We chose objects (and our recruitment strategy – see below) to create variation in the degree of information asymmetry between buyers and sellers. Two of the objects were new, popular tablet computers: the “New Apple iPad” and the “Microsoft Surface”. These two products have a number of features that create rich context for discussion, and since many of these features are likely not known to potential buyers they also create uncertainty about the products’ value. Moreover, the fact that a meaningful proportion of our subject pool already owned one of these two products enabled us to recruit a cohort of sellers with more experience using them than the typical buyer. The other two objects were completely-defined abstract goods: a \$500 check post-dated for January 1st, 2014, and a lottery that paid \$1,000 with 50% chance and \$0 with 50% chance. In contrast to the tablets, there was no uncertainty about the objective attributes of these goods. Buyer valuations for these should in theory depend only on time and risk preferences and hence should be difficult for sellers to influence, at least through the self-regard channel. The order in which subjects discussed these four products was randomized at the session level, and subjects did not learn which product they would be discussing in each conversation until after buyer-seller pairings had been made.

Buyers and sellers were recruited to different buildings on campus. They were not made aware of the other room full of participants until the instructions were read. Buyers sat in the Economics experimental laboratory. Their instructions were read by subjects on the computer (using the Qualtrics survey platform) with only clarifying questions answered. Their valuation elicitation were done in private at their computer terminal once they confirmed the instructions were clear. They then dialed into calls with sellers using cordless landlines. Sellers

sat in individual music practice rooms, preventing them from overhearing the conversations other sellers were having and (potentially) learning from their approaches. Sellers read their instructions off of hard copy slide shows, with only clarifying questions answered. They then dialed into calls with buyers using their cell phones and headsets provided by the researchers. Calls were made and recorded using a third party conference call service.⁷ Calls would end at the seller’s discretion or after ten minutes had passed.

After each conversation and after the ensuing WTP elicitation, buyers filled out an unincentivized survey about the conversation they had just had. For example, buyers wrote down as many factual statements the seller made as they could recall (up to ten) and noted for each whether that statement would be more likely to make someone pay more, less, or about the same for the product, and whether it was more likely to be true or false. The question we make most use of in the analysis below asked buyers to rate how similar the seller was to them on a scale from 0-10.⁸

2.1 Expert Predictions and Conversation Coding

Besides valuations and survey responses, we also viewed the communications between buyers and sellers themselves as data that potentially shed light on how persuasion works. Recent analyses of such content have demonstrated surprising predictive power and insights into the nature of communication (Pennebaker, 2011). We therefore recorded and later transcribed all buyer-seller conversations.⁹

In coding these transcripts for econometric use we were cognizant of two common concerns with content analysis. First, researchers often worry that coding is inherently subjective. We agree, but do not view this as a (necessary) weakness; rather, we take the view that because communication is inherently at least somewhat subjective it is important to capture the “right” subjective interpretation, i.e. the one that is relevant for the research question posed. To do this we recruited coders from the same subject pool as the buyers from our original experiment. This ensures that our coders interpretation of messages are (statistically) the same as those of the audience for which they were originally intended. We also incentivized thoughtful, honest answers by our coders using the coordination game structure proposed by Houser and Xiao (2011): two subjects were anonymously paired with each other, read the same transcript and were asked the same question, e.g. “How likeable did you find this seller?” They were paid every time their answers matched.

⁷rondee.com

⁸The other questions included were: What credible sources did the seller cite and how often? Did they seller make any statements of intent? Did they seller make any promises? (Charness and Dufwenberg, 2006) How surprised would your friends be if they knew how much/little you paid for this product? What percent of OSU students own the product? How likeable was the seller? How trustworthy was the seller? What did the seller think of you? Did the seller listen well? Was the interaction adversarial/cooperative? Did you know the seller beforehand? How clear were the instructions? How confusing was the study? Could anything be improved? Were you given enough time? What was the point of the conversation? And what was the other participant’s strategy?

⁹Due to technical difficulties with the teleconferencing service, the phones themselves, and the online recording and hosting service, we either failed to record or lost the recordings for seven conversations, leaving 257 conversations for the content analysis.

Second, researchers often worry that rich textual data-sets invite data-mining. We therefore sought to discipline our analysis by committing ex ante to a fixed set of informed, theory-relevant predictions. Rather than selecting these ourselves from the literature and attempting to map them into our context, we solicited proposals directly from leading researchers in the field using a novel competitive mechanism. We contacted eight experts from across economics and psychology and asked them to describe five variables each that could be obtained from transcripts using either computer or human coding and that would best predict changes in buyer valuations in our experiment.¹⁰ We provided each expert with the instructions for our experiment and informed them that we would estimate regression models interacting their predictors with product fixed effects, identify the submission with the highest R^2 , and publish the name of the winner (but keep the names of other participants anonymous). This scheme was meant to make the challenge fun and to elicit thoughtful predictions; identifying the “winner” is not central to our results. Seven of the eight experts we invited agreed to participate in the prediction challenge; the full set of variables they recommended we code is listed in Section B. The winner was Joel Sobel, whose variables predicted 18% of the variation in our data.

2.2 Placebo Condition

Unlike many experimental designs in which the obvious counterfactual to “treatment” is “no treatment,” there are several outcomes to which we could meaningfully compare post-conversation valuations. If we want to compare these choices to decisions made after not having given any previous thought to the problem, then pre-conversation valuations are arguably the best counterfactual. At the same time it seems plausible that part of the effect of a conversation is simply the effect of having additional time to consider the decision. In this case it is useful to benchmark buyers’ valuations after conversations with valuations after they have additional time to reflect.

To estimate this benchmark we also ran a placebo condition that was structurally identical to the one above except without conversations: a buyer was matched with a seller who could potentially earn a commission, but they were not allowed to communicate. These buyers instead waited silently for about the length of a phone call before submitting ex post valuation. We ran this placebo treatment during the last three sessions.

3 Recruitment and Descriptive Statistics

We recruited subjects via an email to the entire Ohio State Economics experiment subject pool inviting them to take a survey, which would qualify them for our experiment. This two-minute “intake survey” asked a variety of questions about ownership, use, and familiarity of various products. In order to identify subjects with more or less experience with the tablet goods in our experiment, we asked survey respondents (among many filler questions) about

¹⁰We do not claim, nor did we intend, for the eight expert to be the “top” eight experts in the field of persuasion. Rather, they are all well published and recognized in the field.

their ownership of and familiarity with the iPad (any version) and the Surface.¹¹ We then assigned anyone who filled out the intake survey and said they owned the iPad *or* the Surface to be sellers, and assigned anyone who participated in the intake survey and said they owned *neither* the iPad or the Surface to be a buyer. We did not tell subjects what role they would be playing, so that the only perceptible difference was the location on campus to which subjects were invited.

We held 15 experimental sessions between January and April of 2013, during which we observed 264 buyer-seller conversations and 68 placebo interactions. Appendix Table A-2 summarizes participation by session. In sessions 1 through 11, equal numbers of buyers and sellers participated and all buyers spoke to sellers four times each, once per product. In sessions 12 through 15 more sellers than buyers participated and buyer-product pairs were randomly assigned to one of two conditions: conversations with a seller, as in sessions 1 through 11, or a placebo condition in which the buyer did nothing for ten minutes. The analysis that follows refers to data from conversations in the first 11 sessions except where otherwise noted.¹² Overall, sixty-six sellers participated in four conversations each. Fifty-five buyers participated in sessions in which they spoke with four sellers, while 28 buyers participated in sessions where, randomly from product-to-product, they might have conversations, or they might be in the placebo condition.¹³ Subjects were paid \$15 for participation plus their payments from buyers' choices and the lottery the following week. No session lasted longer than two hours.

Table A-1 in the appendix summarizes participant characteristics. Buyers and sellers are similar demographically but differ, as expected, in their product ownership. These differences are most stark for the iPad and Microsoft Surface as – by design – no buyer owned either of these devices while all sellers owned at least one. Buyers are not entirely unfamiliar with these products, however, as 86% and 16%, respectively, indicate they have previously used one. Differences for other consumer electronics are less stark but point in the same direction. Sellers also spend more than buyers in a typical week.

Figure A-2 displays average initial valuations for the four products. The light gray region shows how much of the total valuation is for the product alone, and the dark gray is the increase in their valuation when the seller receives a commission. Over all of the goods for sale, initial valuations for the product alone is \$279 (significantly different than zero, $p < 0.001$), which increases by \$14 when a commission is added significantly larger than zero ($p = 0.06$). Buyers exhibit a high degree of impatience, valuing a \$500 check nine to twelve months from now \$390. As one would expect, given that they have chosen not to purchase tablets, buyer valuations of the tablets are well below their \$500 market prices at the time. On average, buyers are

¹¹In addition, the intake survey only asked about ownership, use, and familiarity of cars, bicycles, iPhones, iPods, Android tablets, e-book readers (eg Kindle), Windows 7 phones, Android smart phones, paper shredders, scanner/copiers, digital cameras, food processors, and air purifiers in addition to whether they voted in the 2012 presidential election (with one answer identifying non-American citizenship), the number of Economics experiments they have participated in, the number of Psychology experiments they have participated in, what their major is, what their weekly consumption is (in \$), their age, and their gender.

¹²We focus on the first 11 sessions as there could potentially be an interaction between the treatment and control in the latter sessions. The results below do not qualitatively change when conversation data from sessions 12-15 are included (available on request).

¹³Our sample size was not limited by an ex ante power calculation but rather by how many tablet owners we were able to recruit to take our intake and subsequently to participate in the experiment.

willing to pay \$282 for the iPad and \$236 for the Surface. Finally, buyers are quite risk averse, valuing a 50/50 chance of winning \$1,000 at \$211.

4 Results

We lay out methodology and then analyze the data in four stages. First, we examine whether communication changes valuations and whether these changes tend to benefit sellers, in spite of their known conflict of interest. Second, we assess the relative importance of appeals to self-interest versus other-regard. Third, we investigate the roles of sellers strategies, of product characteristics, and of buyer-seller homophily in determining the extent and mechanics of persuasion. Finally, we analyze how persuasion and its pathways are predicted by the identity of the participants.

4.1 Methods

Recall that for the sake of analysis we define subjects' valuations for objects as the highest value X for which they indicated that they preferred to receive the object and $\$500 - X$ to receiving $\$500$. Because we did not impose monotonicity on these responses it was possible for subjects to buy at one price X but refuse to buy at a lower price $X' < X$. In practice, of the 1,328 valuations we elicited, 1,306 (all but 2%) are monotonic in the expected direction.¹⁴

We define the change in buyers' valuation for a product under the commission condition as "persuasion". Let $WTPC_{bsp}^1$ be buyer b 's WTP for product p under the commission condition after a conversation (or waiting period) with seller s , and $WTPC_{bsp}^0$ be her pre-conversation (or pre-waiting period) WTP. Let $WTPF_{bsp}^1$ be the analogous valuations in the fixed (non-commission) condition. Persuasion is then $\Delta WTPC_{bsp} = WTPC_{bsp}^1 - WTPC_{bsp}^0$. To see how this contains changes due both to self-interest and to other-regard, note that we can write

$$\underbrace{\Delta WTPC_{bsp}}_{\text{Persuasion}} = \underbrace{\Delta WTPF_{bsp}}_{\text{Change in Self-interest}} + \underbrace{\Delta (WTPC_{bsp} - WTPF_{bsp})}_{\text{Change in Other-regard}} \quad (1)$$


Note that this is an accounting identity: it says that any change in overall valuation must (mechanically) equal the sum of changes in that valuation's two components. It does *not* rule out the possibility of complex interactions between self-interest and other-regard. For example, manipulating other-regard might be an effective strategy largely because it builds trust which then enables sellers to persuade the buyer to increase the self-interested component $WTPF_{bsp}$ of their valuations.

Where possible we operationalize hypothesis tests non-parametrically using clustered rank-sum tests. Where we need to condition on additional variables to improve precision or remove bias we use regression models that extend the the following baseline specification:

$$Y_{bsp}^1 = \alpha + \beta Y_{bsp}^0 + \gamma_p + \epsilon_{bsp} \quad (2)$$

¹⁴2% is a low number relative to other experiments. Holt and Laury (2002), for example, find non-monotonicity rates of 5.5% and 6.6% in their high- and low-stakes risk preference treatments.

Figure 1: Changes in valuation driven by communication



Histogram of absolute valuation changes, by treatment.

where b indexes buyers, s indexes sellers, p indexes products, and Y_{bsp}^t is a valuation for a product at time $t \in \{0, 1\}$. We include product fixed effects to account for systematic differences between products. The advantage of this approach relative to moving Y_{bsp}^0 to the left-hand side and treating the change $Y_{bsp}^1 - Y_{bsp}^0$ as the dependent variable is that it does not impose $\beta = 1$, which our data easily reject ($p < 0.001$).

4.2 The Mechanics of Persuasion


Result 1. *Communication affects buyer valuations, with effects through both self-interest and other-regard.*

Our first result is that communication altered buyers' choices. Figure 1 shows the distributions of absolute changes in valuations with commission for both the communication and control treatments. Visually, the control distribution has much less spread in the change in valuations. Indeed, absolute changes of valuations for communication conditions is \$49 (for all 15 sessions), significantly and substantially larger compared to \$24 for the control (clustered rank-sum $p < 0.01$).¹⁵ Further, 56% of buyers change their valuation in the control group and 31% in the in placebo.

The impacts of conversation work through changes in both self-interest and other-regard. Figure 2 illustrates this, plotting the histograms of absolute changes for the two components of persuasion. The mean absolute change in self-interest is \$60 following a conversation compared

¹⁵We are underpowered to perform the analysis within the last three sessions alone.

Figure 2: Changes in self-interest and other-regard driven by communication



Histogram of absolute change in self-interest, by treatment.

Histogram of absolute change in other-regard, by treatment.

to \$24 in the placebo (clustered rank-sum $p < 0.01$), and the mean absolute change in other-regard is \$50 if the buyer and the seller speak compared to \$26 if they do not (clustered rank-sum $p < 0.01$). The evidence indicates that changes in valuations observed in our baseline is largely a result of communication rather than other artefacts of our design (e.g. time delay or asking twice).

To quantify the relative contributions of self-interest and other-regard to variation in persuasion, we ask what proportion of the overall variation is eliminated by conditioning on each component. Specifically, we estimate

$$1 - \frac{\mathbb{E}[V(\Delta WTPC | \Delta WTPF)]}{V(\Delta WTPC)} \quad (3)$$

to measure the proportion of total variance attributable solely to self-interest, and

$$1 - \frac{\mathbb{E}[V(\Delta WTPC | \Delta(WTPC - WTPF))]}{V(\Delta WTPC)} \quad (4)$$

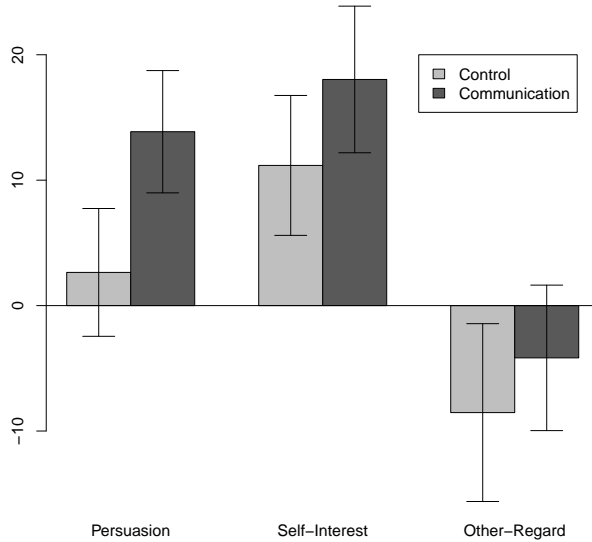
to measure that attributable solely to other-regard. We find that conditioning on self-interest reduces variance more than conditioning on other-regard, but the difference is not dramatic: 16% of the variance can be explained by other-regard alone, and 29% can be explained by self-interest alone.¹⁶

Result 2. *Sellers benefit from communication, and on average this is driven by changes in buyers (perceived) self-interest.*

Communication affects not only the distribution of buyers' willingness to pay but also its mean. In the full 15-session sample, the average change in buyers valuation with commissions following communication is \$14 (all 15 sessions), significantly different from the mean change in placebo of \$3 (clustered rank-sum $p = 0.01$) (See left-most bars in Figure 3). This overall

¹⁶Note that simpler variance decomposition procedures are hard to interpret meaningfully here given the non-independence of the two components.

Figure 3: Communication increases valuations; persuasion works primarily through self-interest



Plots the mean change in self-interest and other-regarding valuation components as defined in Equation 1 as well as total persuasion, for both the placebo and communication treatments. Vertical bars depict standard errors allowing for clustering by buyer and by seller.

difference reflects smaller differences in both self-interest and other-regard, both of which are directionally positive but statistically insignificant. The average change of perceived self-interest is \$18 for the treatment and \$11 for the placebo (clustered rank-sum $p = 0.2$) (See middle bars in Figure 3). The average change in other-regard was -\$4 in the communication arm compared to -\$8 in the placebo arm (clustered rank-sum $p = 0.2$) (See right-most bars in Figure 3). Though insignificant for the two channels, despite the short time frame and known misaligned incentives, sellers systematically increased how much buyers were willing to pay for the goods. Given this, we focus henceforth on changes under the communication treatment and (conservatively) on data from the first 11 sessions with only communication interactions.

Which component of persuasion explains sellers' average gains? We find that mean effects are driven entirely by self-interest. Figure 3 shows that in the full 15-session sample they are driven primarily by self-interest. The same is true in the restricted sample: the mean change in self-interest is \$20 and significantly different from zero (clustered t-test $p < 0.01$), comprising 115% of total persuasion, while the average change in other-regard is -\$2.64 and insignificantly different than zero (clustered t-test $p = 0.7$). In sum, interacting with sellers affects both buyers self-interest and their other-regard, but only the former effect is mean positive in the full sample of participants.

Result 3. *A minority of sellers plan to target other-regard and outperform those who do not.*

One possible interpretation of the results thus far is that while other-regard does affect valuations, sellers have no ability (or intent) to strategically manipulate it. Instead, buyers'

Table 3: Sellers who target other-regard outperform

	Persuasion	Self-Interest	Other-regard	Persuasion	Self-Interest	Other-regard
Targets Other-Regard	25.83 (13.21)*	-1.38 (14.60)	27.21 (13.80)**	34.46 (20.55)*	-8.58 (23.78)	43.05 (23.56)*
Targets Self-Interest	1.81 (17.46)	-5.78 (20.01)	7.59 (17.46)	9.42 (15.00)	-10.30 (21.48)	19.72 (20.99)
Has Any Strategy	11.73 (19.58)	14.10 (20.28)	-2.37 (15.22)	-0.26 (20.43)	1.14 (27.20)	-1.40 (21.41)
Seller FEs	No	No	No	Yes	Yes	Yes
N	220	220	220	220	220	220
R^2	0.04	0.03	0.02	0.22	0.25	0.27

Each column is a separate regression. Each regression includes product fixed effects. “Targets other-regard”, “Targets self-interest”, and “Any strategy” are indicators equal to one if at least two of three coders agreed that the seller’s self-reported strategy involves manipulating other-regard, manipulating self-interest, or influencing the buyer in any way respectively. Standard errors clustered by buyer and by seller in parenthesis. Statistical significance is denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

other-regard adjust upwards or downwards to fixed characteristics of sellers which they discover through conversation.

To test this hypothesis we examine the strategies sellers report to us before each call and their relationship with persuasive performance. Some sellers do unambiguously plan to target the buyer’s self-interest. For example, sellers planned to “talk about how awesome the iPad is,” “tell them the specs and convince them the Surface is worth full retail price,” “give extra information,” and “sell benefits – USB, keyboard for notes, fast internet, touch screen capabilities, no stylus required to take notes.” Yet other sellers describe strategies that explicitly target other-regard. In this group are sellers who planned to “get the customer on my side,” to “build a rapport,” and to “get on buyer’s good side.” Some sellers even articulated specific other-regarding tactics from the literature, as for example the seller who planned to “remember the person’s name! People love when others remember their names.” Motivated by these qualitative comments, we coded seller strategies to identify those that target other-regard, those that target self-interest, and those that seem to contain any intention of influencing the buyer. Three undergraduate research assistants independently assessed whether each seller strategy referred to other-regard. The data they coded was stripped down to subject identifier and seller strategy. We code a seller strategy if at least two coders identified it as such; 19% of strategies were coded as targeting other-regard, 30% were coded for self-interest, and 55% were coded for having any strategy.

We analyze the efficacy of each strategy in Table 3. The third column shows that sellers who use strategies that target other-regard are successful in their stated aim, raising buyers’ other-regard by significantly more than their peers. Interestingly, the second column shows that this does not come at a penalty with respect to their impact on buyers’ self-interest: the difference is small and insignificant. The first column shows the upshot of these results: sellers who target other-regard do significantly better than their peers overall. Moreover, the estimated effect size represents a sizeable 149% of average persuasion in our conversations. Apparently a sizeable minority of sellers do exert influence through other-regard, and these

sellers also do better than average.

This result is important in its own right as it suggests that other-regard may be important for understanding the careers of workers who go on to specialize in persuasion-heavy tasks such as sales or negotiation. Of course, it also raises the question whether these (or other) sellers would have performed better or worse had they used different strategies. While we cannot answer this question conclusively without experimentally manipulating strategies, we can provide some indication by exploiting the fact that some of our sellers changed their strategies between interactions. The last three columns of Table 3 report the results of specifications that include seller effects and thus estimate the within-seller effect of strategy choice. While the estimates are slightly less precise, they are qualitatively similar and if anything more pronounced than the pooled estimates. In particular, sellers have a larger effect on other-regard when using strategies that target other-regard and do better overall as a result. Note, moreover, that these specifications include product fixed effects, so that the result cannot be explained by variation in the underlying decision problem. It thus appears that other-regarding strategies are more successful than alternatives for the subset of sellers who use them.

Result 4. *The nature of the good being sold affects both the extent and the mechanics of persuasion.*


An important hypothesis embedded in our design was that the relative importance of self-interest and other-regard might depend not only on seller strategy but also on the nature of the underlying decision problem, and in particular on the degree of information asymmetry between buyer and seller. Intuitively, asymmetry information would seem to create greater scope for sellers to manipulate perceived self-interest. To examine this we compare outcomes for the two tablet computers of which every seller owned at least one and buyers owned neither, with outcomes for the lottery and time preference questions (“abstract goods”).

We find that sellers were noticeably more successful overall at persuading buyers to pay more for tablets. Figure 4 breaks down average valuation changes, overall and for the two pathways, for the tablets and the abstract objects. The left-most bars indicate that sellers are more successful overall at persuading buyers with the tablets. Valuations increase by \$9.6 for the abstract goods as opposed \$25.1 for the tablets (rank-sum, clustered at buyer and seller level, $p = 0.03$). The middle bars explain why: Changes in self-interest for the abstract goods is \$6 versus \$34 for the tablet computers (rank-sum, clustered at buyer and seller level, $p = 0.05$). Directionally, this came at a cost as other-regard decreases by \$8.9 for the tablets while it increases by \$3.6 for the abstract goods, but this difference is not statistically significant (rank-sum, clustered at buyer and seller level, $p = 0.22$). We interpret these results to mean that the relative importance of other-regard and self-interest in persuasion is conditional on the nature of the decision facing the buyer. They also suggest that self-interest plays a larger role when information is more asymmetric, something that could be tested experimentally by directly manipulating seller information sets.

Result 5. *Persuasion increases in buyer-seller homophily, and in particular gender match.*

We close this subsection by examining the role of homophily, or buyer-seller similarity, in persuasion. We measure similarity in two ways. First, we use buyers’ unincentivized self-

Figure 4: Valuations increase more for consumer electronics; driven by appeals to self-interest




Plots the mean change in self-interest and other-regarding valuation components as defined in Equation 1 as well as total persuasion, for tablets and abstract goods. Vertical bars depict standard errors allowing for clustering by buyer and by seller.

reports of their beliefs regarding how similar the seller was to themselves. The benefit of this measure is that it is comprehensive; the drawback is that buyers might feel a need to answer in a way that rationalized their valuation responses. The evidence generally points towards subjects reporting honestly: report are distributed fairly uniformly across the eleven bins from 1 (“not at all”) to 11 (“very much so”), with no score occurring less than 5% of the time or any score occurring more than 14% of the time. The average is 5.8, and the median is 6. Buyers who were gender-matched with their seller reported a significantly higher similarity score, 1.0 point higher (rank-sum $p < 0.01$). Similarly, buyers matched on nationality with their seller reported they were more similar, 1.5 points higher (rank-sum $p < 0.01$). This suggests the buyers were reporting thoughtfully. As a robustness check, however, we also test directly whether match on the observable characteristics of gender and nationality predict persuasion.

We find that homophily strongly predicts persuasion. Figure 5 illustrates this graphically for self-reported similarity, while Table 4 reports regression coefficients for all three measures. In Column 1 we estimate a \$5.7 increase in persuasion for every one point increase in similarity score ($p < 0.01$). To provide some perspective, this implies that a one standard deviation increase in similarity would produce an \$17.7 increase in persuasion, more than doubling average persuasion (\$17.4). This result is even stronger when we focus on match on gender specifically: while neither buyer or seller gender predicts persuasion or changes in self-interest or other-regard, the interaction of their gender does. The fourth column shows that pairs that are matched on gender have increases of \$22 more than the gender-mismatched pairs

Figure 5: Perceived similarity predicts influence



Plots the mean change in persuasion by how similar the seller was to the buyer according to the buyer, where 11 is most similar. Vertical bars depict standard errors allowing for clustering by buyer and by seller.

($p = 0.04$). Only for nationality is this pattern weaker; persuasion is higher by \$45.4 for same-nationality pairs but this is not statistically significant.

It seems reasonable that the homophily effect could be driven by either channel of persuasion: similarity could beget altruism, build trust, or given the seller insight into the buyer’s preferences, for example. We find evidence for both channels, but somewhat stronger evidence for self-interest. A one point increase in reported similarity increases changes in self-interest by \$4.2 ($p = 0.06$) and other-regard by \$3.0 ($p = 0.04$). With gender-matched pairs this gap is starker: gender-matched pairs see a large \$29.4 increase in valuations via self-interest ($p = 0.02$) but only an insignificant \$3.3 increase via other-regard ($p = 0.6$). Neither channel is significant for nationality match.

4.3 Predicting Persuasion: Buyer and Seller Identities

Result 6. *Persuasion is significantly predicted by who is buying, not by who is selling.*

The results in the previous sub-section describe persuasion among the participant pool at large. This frame is most relevant for understanding applications in which broad swathes of a population are motivated to persuade their peers: for example, viral marketing or the spread of a network good. For other applications, however, it seems more relevant to understand what the *most* persuasive individuals do. Professional salespeople or negotiators, for example, are likely selected for in part because they are relatively good at persuasion.

While we can always test specific hypothesis – such as whether specific seller strategies are linked with success – the panel structure of our data also allows us to conduct a less structured analysis. Specifically, we can estimate the distribution of seller (and buyer) fixed effects and then ask (a) how much better the best sellers are than their peers, and (b) what traits or behaviors correlated with being good at persuasion. The tradeoff is that by virtue of its flexibility this approach is also less powerful statistically, as it is an omnibus test and thus weak against specific alternative hypotheses. For example, if a small minority of sellers have a particular trait that makes them especially influential we will not be well-powered to detect this.

With this in mind, the results summarized in Figure 6 are stark and (to us) surprising. Each column in the figure represents the adjusted R^2 from a regression of persuasion ($WTPC_{bsp}^1 - WTPC_{bsp}^0$) on a particular battery of fixed effects. Recall that the “adjustment” here lets us compare the explanatory power of a given set of variables to that of an equal number of randomly generated variables, and thus allows us to make meaningful comparisons across models with different numbers of predictors. We find that order and product effects are, unsurprisingly, unimportant for explaining the *change* in valuations. Seller effects are even less predictive, performing worse than an equally sized random set of covariates (adjusted $R^2 = -0.07$). They are all noticeably less predictive than buyer effects (adjusted $R^2 = 0.07$). In fact, only the prediction challenge winner was able to explain more variance than buyer fixed effects alone.

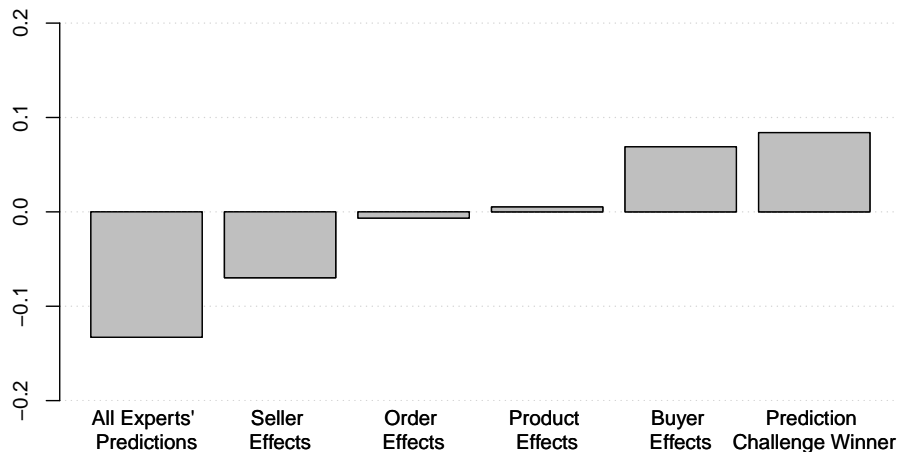
Figure A-1 in the data appendix show the distributions of seller (panel a) and buyer (panel b) fixed effects. Other than three outliers in amongst the sellers, the distribution appears flat relative to the buyer distribution, which is more steadily increasing from left to right. To

Table 4: Buyer-Seller homophily predicts persuasion, change in self-interest, and change in other-regard

	Persuasion	Self-Interest	Other-Regard	Persuasion	Self-Interest	Other-Regard	Persuasion	Self-Interest	Other-Regard
Similarity Score	5.71 (1.53)***	4.22 (2.23)*	3.00 (1.41)**						
Gender Match				22.01 (10.33)**	29.36 (12.41)**	3.28 (6.68)			
Buyer Female				1.53 (12.46)	-5.16 (13.56)	-3.46 (8.79)			
Seller Female				-1.26 (10.29)	-2.54 (13.02)	-14.90 (7.68)*			
Citizenship Match							5.39 (7.12)	0.44 (20.42)	13.45 (17.09)
Buyer Non-American							15.69 (9.46)*	1.13 (16.33)	19.79 (19.97)
Seller Non-American							-11.61 (7.42)	-30.16 (21.93)	16.91 (15.16)
N	210	210	210	220	220	220	216	216	216
R ²	0.74	0.61	0.08	0.74	0.62	0.08	0.74	0.63	0.08

Each column is a separate regression. Each regression includes product fixed effects. Standard errors clustered by buyer and by seller in parenthesis. Statistical significance is denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 6: Buyer identity predicts influence better than seller identity



Plots the adjusted R-squared for OLS regressions, predicting post-conversation commission valuation using initial valuation and covariate(s) listed. Expert predictions for “All Predictions” and “Prediction Challenge Winner” interacted with product fixed effects.

make inferences of these effects, we estimate models that begin with the base specification in equation 2, augmenting it with various sets of fixed effects. For each regression we conduct an F -test of the restriction that the included fixed effects are jointly zero. While we cannot reject the null that seller effects are jointly zero ($p = 0.87$), we easily reject the null that buyer effects are zero ($p < 0.01$).¹⁷ Figure A-1 plots the distributions of these individual effects for reference. The data do not make it clear whether the buyer’s differ in propensity to be persuaded to like a product or a person. Utilizing the same F -test to test for buyer effects on change in self-interest is insignificant ($p = 0.19$) as is the same test for other-regard ($p = 0.53$) (Seller F -test for self-interest $p = 0.64$, other-regard $p = 0.38$). In short, the evidence for *persuasive* types is much weaker than the evidence for *persuadable* types.

5 Conclusion

In this paper we define and apply an experimental framework for studying the mechanics of persuasion: in particular, whether and when persuasion works through appeals to self-interest

¹⁷Standard asymptotic inference for this statistic requires that observations per participant approach infinity. In light of our short panel, we also calculated alternative p -values via randomization inference. Under the null hypothesis that changes in valuation are independent of the individuals involved, the distribution of our data is unaffected by permuting individual identifiers. We therefore draw 1,000 permutations of individual identifiers, re-estimate the F -statistic after each, and use the resulting distribution to estimate the sampling distribution of our original F -statistic (Lehmann and Romano, 2005). As before we cannot reject the null that valuation changes are independent of seller identifies ($p = 0.87$) but reject the null that they are independent of buyer identifies ($p < 0.01$).

or to other-regard. We find that both mechanisms matter to a degree, with substantial variation depending on context and individuals. The “average” seller in our experiment benefits from communication primarily because communication convinces buyers that it is in their own best interest to pay more for objects. At the same time, a minority of sellers make a strategic choice to focus on getting buyers to like them, and these sellers outperform their peers significantly both at cultivating other-regard and at persuasion overall. Sellers gains are driven more by buyer self-interest when they discuss tangible products about which sellers were better-informed, but more by other-regard in settings of symmetric information. Finally (and surprisingly), variation in persuasion is driven less by who is selling than by who is buying.

Our experimental framework can be adapted to study a wide range of additional questions about persuasion. We close by highlighting three. First, future work could compare the effects on buyers of conversations with expert as opposed to amateur persuaders. Many roles in an economy require specialized persuasive skills – salespeople, CEOs politicians, and so forth – and it would be valuable to understand both what it is that sets such expert persuaders apart and where these skills come from – whether innate or acquired.

Second, future work could compare the effectiveness of persuasion with and without face-to-face interaction. Firms spend huge sums of money flying their employees (and especially their salespeople) around the world in the belief that face-to-face interactions yield higher returns than telephone or even videoconferenced interactions, yet there is little rigorous evidence on those returns. Such experiments could also shed light on the evolution of urban form; as scholars have noted, rapid ongoing urbanization seems hard to reconcile with dramatic improvements in communication technology unless face-to-face interactions are particularly valuable.

Third, future work could compare persuasion about private goods to persuasion about public ones. Consider for example persuading a neighbor to vote in favor of environmental reforms: this problem has the interesting features that (i) it may never be in the neighbors narrow self-interest to do so, and (ii) voting for the reform likely benefits the persuader if and only if it also benefits the neighbor, creating a direct link between self-interest and other-regard that is absent in the private goods case.

References

- Andreoni, James and Justin M. Rao**, “The power of asking: How communication affects selfishness, empathy, and altruism,” *Journal of Public Economics*, August 2011, *95* (7-8), 513–520.
- Bohnet, Iris and Bruno S. Frey**, “The sound of silence in prisoner’s dilemma and dictator games,” *Journal of Economic Behavior and Organization*, 1999, *38*, 43–57.
- Carnegie, Dale**, *How to Win Friends and Influence People*, Simon and Schuster, 1936.
- Charness, Gary and Martin Dufwenberg**, “Promises and Partnership,” *Econometrica*, November 2006, *74* (6), 1579–1601.
- and **Uri Gneezy**, “What’s in a name? Anonymity and Social Distance in Dictator and Ultimatum Games,” *Journal of Economic Behavior and Organization*, 2008, *68*, 29–35.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik**, “Pandering to Persuade,” *American Economic Review*, February 2013, *103* (1), 47–79.
- Cialdini, Robert**, *Influence: the Psychology of Persuasion*, William Morrow and Company, Inc., 1993.
- Crawford, Vincent P and Joel Sobel**, “Strategic Information Transmission,” *Econometrica*, November 1982, *50* (6), 1431–51.
- Currarini, Sergio, Matthew O. Jackson, and Paolo Pin**, “An Economic Model of Friendship: Homophily, Minorities, and Segregation,” *Econometrica*, 07 2009, *77* (4), 1003–1045.
- DelaVigna, Stefano and Matthew Gentzkow**, “Persuasion: Empirical Evidence,” *National Bureau of Economic Research Working Paper No. 15298*, 2009.
- Dziuda, Wioletta**, “Strategic argumentation,” *Journal of Economic Theory*, July 2011, *146* (4), 1362–1397.
- Feigenberg, Benjamin, Erica M. Field, and Rohini Pande**, “Building Social Capital Through MicroFinance,” NBER Working Papers 16018, National Bureau of Economic Research, Inc May 2010.
- Glaeser, Edward L., David I. Laibson, Jos A. Scheinkman, and Christine L. Soutter**, “Measuring Trust,” *The Quarterly Journal of Economics*, August 2000, *115* (3), 811–846.
- , **David Laibson, and Bruce Sacerdote**, “An Economic Approach to Social Capital,” *Economic Journal*, November 2002, *112* (483), 437–458.
- Golub, Benjamin and Matthew O. Jackson**, “How Homophily Affects the Speed of Learning and Best-Response Dynamics,” *The Quarterly Journal of Economics*, 2012, *127* (3), 1287–1338.
- Greiner, Ben, Werner Guth, and Ro’i Zultan**, “Social communication and discrimination: a video experiment,” *Experimental Economics*, 2012, *15* (3), 398–417.
- Holt, Charles A. and Susan K. Laury**, “Risk Aversion and Incentive Effects,” *American Economic Review*, December 2002, *92* (5), 1644–1655.

- Houser, Daniel and Erte Xiao**, “Classification of Natural Language Messages Using a Coordination Game,” *Experimental Economics*, 2011, 14 (1), 1–14.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, October 2011, 101 (6), 2590–2615.
- Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp**, “Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment,” *Quarterly Journal of Economics*, 2006, 121 (2), 747–782.
- Lehmann, Erich and Joseph Romano**, *Testing Statistical Hypotheses*, Springer, 2005.
- Leider, Stephen, Markus Mobius, Tanya Rosenblatt, and QA Do**, “Directed altruism and enforced reciprocity in social networks,” *Quarterly Journal of Economics*, 2009, 124 (4), 1815–1851.
- Loewenstein, George, Daylian M. Cain, and Sunita Sah**, “The Limits of Transparency: Pitfalls and Potential of Disclosing Conflicts of Interest,” *American Economic Review*, May 2011, 101 (3), 423–28.
- Marmaros, David and Bruce Sacerdote**, “How Do Friendships Form?,” *The Quarterly Journal of Economics*, 02 2006, 121 (1), 79–119.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer**, “Coarse Thinking and Persuasion,” *The Quarterly Journal of Economics*, 05 2008, 123 (2), 577–619.
- Olszewski, Wojciech**, “Informal communication,” *Journal of Economic Theory*, August 2004, 117 (2), 180–200.
- Pennebaker, James W.**, *The Secret Life of Pronouns: What our Words Say About Us*, Bloomsbury Press, 2011.
- Petty, Richard E. and John T. Cacioppo**, *Communication and Persuasion: central and peripheral routes to attitude change*, Springer, 1986.
- Shue, Kelly**, “Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers,” *Review of Financial Studies*, 2013, 26 (6), 1401–1442.
- Small, Deborah A. and George Loewenstein**, “Helping a Victim or Helping the Victim: Altruism and Identifiability,” *Journal of Risk and Uncertainty*, 2003, 26 (1), 5–16.
- Yamamori, Tetsuo, Kazuhiko Kato, Toshiji Kawagoe, and Akihiko Matsui**, “Voice matters in a dictator game,” *Experimental Economics*, December 2008, 11 (4), 336–343.

A Data Appendix

Table A-1: Participant characteristics

Variable	Buyers		Sellers	
	Mean	SD	Mean	SD
Age	21.3	1.9	21.2	1.8
Gender = Male	0.55	0.50	0.55	0.50
Voted	0.57	0.50	0.67	0.47
Owns an iPad	0.00	0.00	0.74	0.44
Used an iPad	0.86	0.35	0.97	0.17
Owns a Windows Tablet	0.00	0.00	0.20	0.43
Used a Windows Tablet	0.16	0.37	0.53	0.50
Owns an iPhone	0.49	0.50	0.71	0.45
Owns an iPod	0.72	0.47	0.76	0.43
Owns an Android Tablet	0.02	0.15	0.11	0.31


Reports means and standard variations of age (top row) and indicator variables for whether a given statement is true (all other rows). By sample construction, no buyers owned an iPad or a Windows table.

Table A-2: Participation by Session


Session	Date	Buyers	Sellers	Conversations	Placebos
1	January 15	4	4	16	0
2	January 22	6	6	24	0
3	January 24	5	5	20	0
4	January 29	5	5	20	0
5	January 30	7	7	28	0
6	February 6	4	4	16	0
7	February 7	5	5	20	0
8	February 14	5	5	20	0
9	March 6	5	5	20	0
10	March 7	5	5	20	0
11	March 28	4	4	16	0
12	April 4	10	6	20	20
13	April 10	7	3	8	20
14	April 17	6	2	4	20
15	April 18	5	4	12	8
Totals		83	70	264	68

Figure A-1: Distribution of participant fixed effects

(a) Sellers




(b) Buyers



Plots the sorted mean change in persuasion individually for each seller and each buyer, respectively.

Figure A-2: Initial valuations show limited altruism



Plots buyers' mean initial valuations, in dollars, for each of the four products studied. Valuations have been decomposed into the mean willingness to pay absent any commission for the seller ("self-regarding") and the mean incremental willingness to pay when commissions are added ("other-regarding").

B Expert Predictions for Correlates of Persuasion

Here, we present the predictions from the persuasion research experts who participated in our *Prediction Challenge*. The goal of the contest was to choose five features of the conversations that would best predict variance in net influence. We present them in no particular order.

Expert 1

1. On a scale of 1 to 5, how much do you think this conversation increased the buyer's valuation for the product? (1=not at all, 5=a lot)
2. On a scale of 1 to 5, how much do you trust this seller? (1=not at all, 5=a lot)
3. On a scale of 1 to 5, how much do you like this seller? (1=not at all, 5=a lot)
4. Is the buyer a man?
5. Is the buyer a man and the seller a woman?

Expert 2

1. The length (in words) of the conversation.
2. The number of times that the seller says "I" (first person singular).
3. The number of times the seller calls the buyer by first name.
4. The total number of times "no" is used.
5. The number of "overlaps/interruptions" (but people talking at the same time).

Expert 3

1. How many statements of the form "I've used it and really like X" does seller make?
2. How many statements of the form "someone I know has used it and really liked X" does the seller make?
3. How many statements of the form "you can use it for X" does the seller make?
4. How many statements of the form "I'm not telling you this because of my incentives" does the seller make?
5. The interaction $(4) \times [(1) + (2) + (3)]$

Expert 4

1. Was the conversation about pros and cons of the object/the decision of buying, or was it tangential?
2. Did the seller appear impartial by giving also cons of buying the object?
3. The duration (in words) of the conversation (though not certain in which direction this will go)
4. The fraction (of words) of the conversation during the buyer spoke.
5. How many questions RELATED to the buying decision (not questions about jobs or weather) did the seller ask?

Expert 5

1. Look for the mere number of words used for buyer and seller
2. How many times does the seller say the buyer's name, the word "you", or the word "your"?
3. On a scale from 1 to 5, how many and how strong do you think the sellers arguments are? (1 = no arguments or only very weak arguments, 5 = many strong arguments)
4. On a scale from 1 to 5, how warm or likeable do you find the participants (1 = not likeable at all, 5 = very likeable)
5. On a scale from 1 to 5, how competent or expert do you find the participants (1 = not at all competent or expert,, 5 = very competent or expert)

Expert 6

1. How many times did the seller lie about or exaggerate the product attributes?
2. Does the seller making a 'panhandler' kind of plea for help?
3. Does the seller describe details about the products (for the two products)?
4. Does the seller provide the actuarial value of the items (for the check and lottery)?
5. Does the seller explain the incentive scheme to the buyer?

Expert 7

1. Word count (the more words in the conversation, but especially the buyer)
2. Language style matching plus positive emotion words (LSM – this uses LIWC and calculates the relative use of the following function words: personal pronouns, impersonal pronouns, articles, prepositions, auxiliary verbs, negations, adverbs, conjunctions). In other words, high LSM and high positive emotion words will result in the highest amount.
3. What did the buyer say they were willing to pay at the beginning of the conversation – or if (s)he did not say, what do you think s(he) was willing to pay at the beginning of the conversation?
4. On a scale from 1 to 5, how personal does the buyer get? (Personal is defined by the buyer saying something about his background, emotions, or personal concerns. Also, more likely to say his/her name.) (1 = not at all personal, 5 = very personal)
5. One a scale from 1 to 5, how would you rate the buyer's initial desire or interest in the product?

Items added by authors

1. Do the buyer and seller make a side deal (agreed to payments outside of the experiment) and/or agree to meet afterwards?
2. On a scale from 1 to 3, how well does the buyer appear to understand the rules of the game? (1 = not at all, 2 = somewhat, 3 = very well)
3. On a scale from 1 to 3, how well does the seller appear to understand the rules of the game? (1 = not at all, 2 = somewhat, 3 = very well)