

Pathways of Persuasion*

Lucas Coffman[†]
Ohio State University

Paul Niehaus[‡]
UC San Diego

January 20, 2016

Abstract

How do people persuade? While economic theories of persuasion emphasize self-interest, others posit a role for other-regard. For example, a salesperson might describe product features but also try to build rapport. We introduce an experimental framework to study these pathways. In our framework, sellers use free-form conversation to convince buyers to raise their valuations for objects. We find sellers benefit from communication despite their conflict of interest. Self-interest explains more variation overall, but a minority of sellers target other-regard and substantially outperform their peers. Results vary across goods, in particular with the degree of information asymmetry, but persuasion is possible even for time and risk preferences. We also observe a trade-off between pathways: sellers' gains along one come at a cost along the other. Across subjects, the "persuadability" of the buyer matters, but the persuasiveness of the seller does not, and buyer-seller homophily (especially gender match) strongly predicts persuasion.

*This research was in part funded by the Ohio State Behavioral Decision Making initiative, the UC San Diego Research Committee, and a UC San Diego Hellman Award. We thank John Beshears, Jason Blevins, Sanford Braver, Katherine Coffman, Stefano DellaVigna, Christine Exley, Paul Healy, John Kagel, Ian Krajbich, Jeffrey Naecker, Muriel Niederle, Andres Santos, Joel Sobel, and seminar participants at UC Berkeley Haas School, Brown University, SITE Experimental, UC San Diego Rady School, UC San Diego, UC San Diego Psychology, University of Pittsburgh, Stanford University, Harvard University, UC Santa Barbara, Cornell, Michigan, the Florence Workshop and ESA Santa Cruz for helpful comments. This paper benefited greatly from the research assistance of William Biscarri, Erica Ng, Haochi Zhang (main experiments), Natalie Mouzooni (coding experiments), and John Conlon (analysis). 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

How do people persuade? A rich theoretical literature has examined when and how persuasion is possible – distinguishing, for example, between information-based persuasion (e.g. Crawford and Sobel (1982), Kamenica and Gentzkow (2011))¹ and taste-based persuasion (e.g. Becker and Murphy (1993)), and more recently between rational and behavioral methods (e.g. Mullainathan et al. (2008), Murphy and Shleifer (2004)).

One common feature of these models is that senders convince receivers that a given action is in their own *self-interest*. For example, a salesperson might describe appealing features of a product, or a friend might share positive reviews of a movie. Popular writing on persuasion, on the other hand, often emphasizes the role of *other-regard*: being likeable, building rapport, and so on. Early editions 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). Much of Robert Cialdini’s popular book *Influence* is devoted to similar arguments, suggesting “people prefer to say yes to individuals they know and like” (Cialdini, 1993). For example, a salesperson might discuss a common interest, or a friend might tell you how much they want to see a movie. This view seems plausible given evidence that other-regard matters in other real-world settings such as workplace productivity (Bellemare and Shearer, 2009), workplace satisfaction (Card et al., 2012), labor relations (Krueger and Mas, 2004), and charitable giving (Andreoni and Payne, 2013).² The literature thus begs a simple but important question: what are the (relative) roles of self-interest and other-regard in persuasion?

We present an experimental framework for studying that question, and results from an initial application. We define an experimental environment in which (1) senders have an incentive to persuade receivers; (2) senders and receivers have flexibility to communicate as they wish (in contrast to experiments which restrict the message space in order to test specific models); (3) we can measure the effects of communication; and (4) we can decompose these effects into components attributable to self-interest and other-regard. We also require that the framework lets us manipulate other aspects of the game (e.g. participants, information sets, the receiver’s decision problem) in order to understand how persuasion varies with context. Per se the framework is intentionally descriptive: it lets us understand what people do when they persuade each other. One can easily use it to test more specific, causal hypothesis (for example, how do results change if we coach sellers to be friendly?), but in this paper we begin

¹See also Olszewski (2004), Dziuda (2011), and Che et al. (2013), among many others.

²Economic analyses of persuasion and of other-regarding preferences have remained largely siloed. As a crude indicator, the four papers on other-regarding preferences with the most Google Scholar citations (Fehr and Schmidt (1999), Rabin (1993), Bolton and Ockenfels (2000), and Charness and Rabin (2002)) are cited by zero economics papers with “persuasion” in the title. There is of course evidence that communication can affect other-regard more broadly. See for example Yamamori et al. (2008), Andreoni and Rao (2011), and Greiner et al. (2012) on communication in dictator games; Charness and Gneezy (2008) and Bohnet and Frey (1999) on (removing) anonymity, Landry et al. (2006) on face-to-face interaction and attractiveness, Roth (1995), Dawes et al. (1977), and Buchan et al. (2006) on irrelevant versus strategic conversation, and Brandts et al. (forthcoming) on rapport-building conversation. There is also related work on how material gifts affect other-regard and subsequent behavior. For example, the gift exchange literature surveyed by Cooper and Kagel (in press) explores how employees reciprocate above-market wages with increased productivity; Malmendier and Schmidt (2012) show that bribes can affect willingness to purchase from the briber.

by simply documenting how persuasion works, absent our intervention.

To make the design concrete, we focus on a buyer-seller paradigm. Subjects assigned to the role of buyers have a telephone call with sellers who receive a commission that increases (in expectation) with the buyer’s willingness to pay for an object, such as an iPad.³ We let them chat however they wish: we do not train or coach the sellers in any way. We capture information about the interaction including seller strategies, buyer perceptions, and the content of the conversations themselves, which we transcribe and code.

We then measure changes in buyer valuations, and decompose these into changes in other-regard and (perceived) self-interest. Specifically, we elicit buyers’ willingness to pay for objects before and after conversations, in two ways each time: when the seller does *not* receive a commission, and when she does. Changes in valuations without a seller commission capture the effects of persuasion on self-interest, while changes in the *difference* between the two measures capture changes in other-regard. Note that this is purely an accounting distinction and does not rule out interactions, such as the possibility that buyers put more faith in advice from sellers they like.⁴ We also benchmark changes in valuation that result from conversations against those that result from giving sellers a similar amount of time to think on their own.

Our framework is flexible enough to vary many contextual factors – the expertise of the participants, the decision problem, the communication channel, the symmetry of information, and so on. We take a first step by varying a few of these. We repeat the experimental procedure four times within each session, rematching buyers and sellers to discuss four different objects for sale. This panel structure lets us estimate individual effects, as well as vary the attributes of the object being sold. We chose objects to vary the degree of information asymmetry between buyers and sellers. Two are consumer electronics devices with which sellers had (by design) more familiarity than buyers. The other two are abstract: a lottery ticket with known payoff distribution and a post-dated check. Because sellers have no private information about these objects, we expected them to be hard to sell by appeals to self-interest.

We first check whether communication does in fact affect buyer behavior. This is not obvious as sellers have a material conflict of interest, so that in a simple “cheap talk” interpretation of our setup they should have no influence. Yet buyers change their valuations by an (absolute) average of 16% following conversations, significantly higher than the average without communication (7%). On the extensive margin, 56% of valuations change after conversations compared to 31% without communication. Communication also benefits sellers: buyers’ mean valuations increase significantly by 5%, compared to an insignificant 1% in no communication. These are large effects, especially given that phone calls lasted just 3.5 minutes on average: extrapolating linearly (which is unrealistic, but entertaining), our amateur sellers generated \$235 of additional buyer product valuation per hour.

What drives these effects? Sellers’ self-described strategies and their behavior in conver-

³Sellers’ incentives are common knowledge, so that disclosure is not a salient choice. See Loewenstein et al. (2011) for evidence on the effects of disclosing conflicts of interest.

⁴The dichotomy between self-interest and other-regard is related to but distinct from that in DellaVigna and Gentzkow (2009) between belief-based and preference-based persuasion. To illustrate, beliefs could affect both perceived self-interest (“it seems like a great product”) and other-regard (“he seems like a great guy”). It is also distinct from dual theories of information processing in psychology (e.g. Petty and Cacioppo (1986)).

sations vary widely. Many focus on the self-interest pathway, planning to “talk about how awesome the iPad is” or “sell benefits,” for example, and making pitches like this:

Seller: I love it cause, it seems great for taking notes and everything like that and it’s very handy cause it’s much lighter than a laptop.

Yet other sellers plan to “build a [sic] rapport,” “get on buyer’s good side,” or “remember the person’s name! People love when others remember their names.” In some cases we even see sellers deliberately sidetrack product conversations to ask about buyers’ personal lives instead:

Seller: That’s good, that’s good, man. Look, are you a student here also?
...What’s your major?

These seller behaviors beg the question whether persuasion works by actuating buyers’ self-interest or other-regard (or both).

On average we find that both pathways matter, but that self-interest accounts for nearly twice as much variation. Other-regard changes in a majority of conversations (52%), with a mean absolute change of 16%, and both figures are significantly higher than without communication (34% and 8%, respectively). Perceived self-interest, meanwhile, changes in 60% of conversations with a mean absolute change of 20% (v.s. 32% and 8% without communication). In a variance-decomposition exercise, self-interest accounts for 29% of the overall variation in persuasion as opposed to 16% for other-regard.⁵ Moreover, changes in other-regard tend to cancel each other out across conversations, so that on average other-regard does not change while private valuations increase significantly.

These results raise the possibility that self-interest is the main pathway for endogenous persuasion, but that changes in other-regard are simply “white noise”. For example, it could be that buyers like some sellers more and others less than expected, but that sellers do not actively influence these perceptions. Yet for a minority of sellers, targeting other regard does seem to be an intentional and effective strategy. In the 19% of cases where we coded seller strategies as targeting other-regard, sellers out-perform their peers by a wide margin, increasing buyer valuations by nearly three times the average, with the bulk of this difference working through other-regard. Even comparing conversations within-seller we find that the same seller does significantly better when targeting other-regard than when not. Manipulating other-regard thus appears to be a highly effective strategy for the minority who use it.

The importance of each pathway also varies across goods. As expected, sellers are more effective at selling electronics than abstract goods. This is consistent with the conjecture that information asymmetry is an important input into persuasion. Yet sellers do manage to significantly increase buyer valuations for the lottery and post-dated check, despite being symmetrically informed. And these gains are differentially due to improvements in buyer other-regard, compared to gains when selling electronics. These results are hard to reconcile with models of fixed, self-interested preferences, and at least suggestive of the idea that other-regard plays a larger role in persuasion when information is less asymmetric.⁶

⁵Note that these figures do not sum to 100% because self-interest and other-regard are not independent, so that some portion of persuasion cannot be attributed exclusively to either. See Section 4.2.

⁶A caveat to the latter result is that we find similar (albeit insignificant) patterns with no communication.

This last result highlights the potential for interaction between the pathways. There could be both mechanical substitution (e.g. if sellers allocate scarce time across strategies) and causal interactions (e.g. if putting on the “hard sell” irritates buyers, or giving helpful advice pleases them). In our data, changes in self-interest and other-regard are significantly negatively correlated with each other. This is true both unconditionally (-0.64) and after projecting the data onto covariates to remove potentially correlated measurement error (-0.70). Moreover, this relationship is roughly constant across quartiles of conversation length, suggesting it is not purely mechanical. Sellers thus appear to face a trade-off, with gains along one pathway partially offset by losses along the other.

In addition to examining the pathways of persuasion, our framework lets us examine the role of buyer and seller identities. We close by highlight two intriguing features of our data. First, persuasion is much better predicted by who is *buying* than who is selling. Specifically, we cannot reject the null that seller fixed effects are identical, while buyer fixed effects are strongly significant and explain more variation than do seller fixed effects, product fixed effects, and order effects combined. Understanding who is *persuasive* can thus be less important in some settings than understanding who is *persuadable*.⁷ Second, homophily (i.e. similarity between buyer and seller) strongly predicts persuasion. A 1 s.d. increase in (perceived) similarity is associated with a 101% increase in persuasion. Interestingly, both other-regard and self-interest increase, suggesting that homophily may enhance credibility as well as directly affecting liking. This is particularly evident for gender similarity: persuasion is 123% higher for gender-matched pairs, due entirely to perceived self-interest.⁸

For research on persuasion, our results raise (we hope) as many questions as they answer. It is an open theoretical question which of the various models of other-regard – e.g. those based on distributional fairness, intentions, reciprocity, or moral perceptions – can best rationalize its role in persuasion (see Cooper and Kagel (in press) for a recent survey). Incorporating other-regard into models of persuasion, and capturing the tradeoffs we see between persuasion strategies, might be illuminating. There are also a number of additional empirical questions which would be straight-forward to test within our framework, and which we discuss in the conclusion – 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)?

Our results also add to work on other-regarding preferences. While well-known and regarded, this literature has made slower headway into real-world applications. We show that other-regard plays a role in persuasion, a central feature of phenomena of economic life (and estimated to account for as much as 25-30% of the US economy (McCloskey and Klamer, 1995; Antioch, 2013)). In doing so we also add to evidence on the determinants of other-regard. A

⁷To our knowledge, little is known about persuadability. Cialdini et al. (1974) find that more easily persuaded confederates were *perceived* as more (less) intelligent by persuaders (observers).

⁸Psychologists have found that people are more likely to report liking (Emswiller et al., 1971) or being willing to do small favors for (Byrne, 1971) those similar to them. Even (deceptively) being told you share a birthday with someone increases willingness to do favors (Burger et al., 2004), cooperate in a prisoner’s dilemma (Miller et al., 1998), and (hypothetically) buy a gym membership (Jiang et al., 2010). Our result is consistent with these findings, but our study differs by measuring revealed as opposed to stated preferences and by decomposing effects into self-interest and other-regard.

number of studies have documented the role of environmental factors;⁹ our paper adds to a smaller literature showing how economic incentives affect social preferences.¹⁰

Finally, our results on homophily contribute to a growing literature documenting its importance for outcomes including trust (Glaeser et al., 2000), social learning (Golub and Jackson, 2012), and network formation (Currarini et al., 2009). The fact that homophily predicts persuasion is particularly important as it could help explain inequitable outcomes for minority groups. Women seeking to advance in professions dominated by men, for example, might find it systematically harder than male colleagues to win support for their ideas.

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 suggests directions for next steps.

2 Experimental Design

In constructing the experimental framework, we had four main objectives. To document broadly how persuasion works, the first goal in designing the experiment was to create relatively natural and relevant settings for persuasion. To increase our predictive power of persuasion, the second goal was to capture rich ancillary data, including conversations themselves. To test the relative importance of the pathways of persuasion, the fourth object is that the experiment enables us to isolate whether persuasion occurred through gains in perceived self-interest or other-regard. Finally, to be useful beyond our own investigation, the fourth goal is to create an environment where numerous factors can be altered – the technology of communication, types of participants, objects for sale, and so on.

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 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 an arm without communication, 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

⁹Such factors include 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)).

¹⁰See for example Glaeser et al. (2002) who document evidence of social capital investment.

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, Panel A.) 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 under the additional condition that the seller will receive a \$500 commission if a sale takes place. 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. (Table 1, Panel B.) 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 \$500.¹¹ All buyers reported their valuations both for the good alone, and for the good with the commission.

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.¹² 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 number of questions in which the buyer selects the alternative in which the seller receives a commission: since 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).

We define persuasion formally as the change in buyers' valuation for a product under the commission condition. Let $WTPC_{bsp}^1$ be buyer b 's WTP for product p under the commission condition after a conversation 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 decompose this into changes due both to self-interest and to other-regard, write

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

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

¹¹We did not size commissions to reflect typical real-world rates, but rather to create meaningful incentives for sellers to exert effort, knowing that any given buyer choice had a low chance of being implemented.

¹²To be able to tell the sellers a precise probability of receiving a commission, we stratified the drawing so that an equal number of commission and non-commission questions would be chosen.

Table 1: Eliciting Willingness to Pay, Flat Wage

Panel A: 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

Panel B: With 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. Buyers answered all questions before proceeding with the experiment. All questions in Panel A were displayed on the same page, with Panel B on the subsequent page.

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. We examine these issues in Section 4.4.

As mentioned in the goals, our experimental design is intentionally flexible, with scope to vary just about any important predictor of persuasion. We are able to take first steps in exploiting this potential. To do so, we introduce some variation in the set of objects for sale and in buyer-seller matches. 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 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, a meaningful proportion of our subject pool already owned one of these two products, which enabled us to recruit a cohort of sellers with more experience using them than the typical buyer (see below). 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 the Ohio State University 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.¹³ 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 unincorporated 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 0-10.¹⁴

2.1 Researcher Predictions and Conversation Coding

Besides valuations and survey responses, messages between buyers and sellers themselves are potentially useful data. Recent work has found that such content can have surprising predictive power and yield insights into the nature of communication (Pennebaker, 2011). We therefore record and transcribe buyer-seller conversations.¹⁵

In coding these transcripts we are cognizant of two common concerns with content analysis.

¹³rondee.com

¹⁴The 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?

¹⁵Due 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.

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 it is important to capture the *relevant* subjective interpretation, i.e. the one that is appropriate to the research question. In our case the question is how messages affected buyers, so we recruit coders from the same subject pool as the buyers from our original experiment. Our coders’ interpretations of messages are thus the same (statistically) as those of the original audience. We also incentivize thoughtful, honest coding using the coordination game structure proposed by Houser and Xiao (2011): two coders are anonymously paired with each other, read the same transcript, answer the same question (e.g. “How likeable did you find this seller?”), and are paid if their answers match.

Second, researchers often worry that rich textual data-sets invite data-mining. We therefore discipline our analysis by committing *ex ante* to a fixed set of metrics. Rather than deriving these ourselves from the literature, we solicited proposals directly from leading researchers in the field using a novel competitive mechanism. We contacted eight researchers 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, that would best predict changes in buyer valuations.¹⁶ We provided each researcher 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, and in doing so to provide us with a snapshot view of how the profession currently thinks about persuasion.¹⁷

2.2 No Communication Condition

There are several possible counterfactuals for buyers’ valuations changing after speaking with sellers. If we want to compare these to what buyers would pay without having given any previous thought to the problem, then pre-conversation valuations are arguably the best counterfactual. If part of the effect of a conversation is simply the effect of having additional time to consider the decision, on the other hand, it is interesting to benchmark against the effects of time to think. To do so we also implement a second condition, identical to the one above but without conversations: we match buyers to sellers, but do not let them communicate. Instead, buyers in this condition wait silently for about the length of a phone call before submitting *ex post* valuations. Any treatment effects in this condition can be attributed to reflection and similar effects. Differences between the treatment effects in the No Communication and the Communication treatment be attributed to the conversations held between participants. Observing effects of a time delay in this condition thus will not affect the interpretation of our results. We ran the no communication condition during the last three sessions.

¹⁶We do not claim, nor did we intend, for the eight researchers to be the “top” eight in the field of persuasion. Rather, they are all well published and recognized in the field.

¹⁷We congratulate the winner, Joel Sobel, whose measures predicted 18% of the variation in persuasion.

3 Recruitment and Data Description

3.1 Recruitment

We recruited subjects via an email to the entire Ohio State Economics experiment subject pool inviting them to take a two-minute intake survey, which would qualify them for our experiment. This survey asked a variety of questions about ownership of, use of, and familiarity with 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 their ownership of and familiarity with the iPad (any version) and the Surface.¹⁸ 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 no communication pairings. Appendix Table A-1 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 buyers than sellers 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 no communication 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.¹⁹ 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 in which they were randomly assigned to the communication condition twice and the no communication condition twice.²⁰ 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.

3.2 Data Description

Table A-2 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

¹⁸In 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.

¹⁹We focus on the first 11 sessions as there could potentially be an interaction between the communication and no communication arms in the latter sessions. The results below do not qualitatively change when conversation data from sessions 12-15 are included (available on request).

²⁰Our 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.

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.

Turning to valuations, we first examine the consistency of the data. Recall that because we did not impose monotonicity on subject responses it was possible for a subject to report that she would pay one price X for an object but refuse to pay 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, a high rate of consistency relative to other experimental studies.²¹

Figure A-1 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 light blue is the increase in their valuation when the seller receives a commission. The mean initial valuation across all products for the good itself is \$279, which increases by \$14 when a commission is added ($p = 0.06$). The latter is not an obvious result a priori: while altruistic behavior is common in laboratory environments like dictator games, we find it here despite having framed a more adversarial buyer-seller interaction.

Buyer valuations of the tablets are generally well below their \$500 market price at the time, with mean valuations at \$282 for the iPad and \$236 for the Surface. This is expected given that buyers have revealed a preference not to purchase at the market price. Buyers are impatient, valuing a \$500 check nine to twelve months from now at \$390, and quite risk averse, valuing a 50/50 chance of winning \$1,000 at \$211.

3.3 Example Conversations

While the full corpus of conversations is too large to present, we provide here a few examples to provide context for the numerical analysis that follows. We selected examples to try to illustrate the range of conversations that took place, and in particular the different approaches that sellers took to selling.

Our sellers were amateurs and, not surprisingly, some capitulated quickly:

Seller: Yeah, if I were--yeah if I were in your position, I would probably think like only get like, 200, \ \$300 profit. That would be a very nice deal already.
Buyer: Yeah, definitely.
Seller: [long pause] Well...[long pause] I don't know. I'm not a good sales person.
Buyer: It's okay.
Seller: We can talk about whatever.
Buyer: So what year are you?
...

Other sellers seem far more comfortable, even to the point of monopolizing the conversation:

Seller: And uh I know judging by your lack of gambling experience and nervousness with that, I think you would probably agree with me there.

²¹For example, Holt and Laury (2002) obtained non-monotonic responses to 5.5% and 6.6% in their high- and low-stakes risk preference elicitations.

But at this point, I would---I would say you might as well go for it. I would do it. Cause you made fifteen bucks here. So, you're positive. You're not going to lose any money. I'd say-I'd say it's not a safe risk, but fifty-fifty shot on winning a \\$1000 in addition to the big one five that you're pulling in right now as your [inaudible]. I don't --I don't--I don't think I could turn that down. I think I would have to go for it.

Buyer: Alright. Well, I'll keep that in mind.

...

Some conversations focus on the attributes and value of the object for sale:

Seller: Haha, I'm pretty good. So, I'd like to talk to you a little bit more about the iPad as far as, do you own an iPad or anything? Er, do you own an iPhone?

Buyer: Uh, I have only, I'm using an apple right now.

Seller: Ok so, the iPad obviously, to me, I personally have an iPad, and I, I love it cause, it seems great for taking notes and everything like that and it's very handy cause it's much lighter than a laptop. So that's kinda nice.

Buyer: Uh, huh.

Seller: Uh, I would definitely recommend getting one. I know they are on the pricier side, but as far as I know they are much better iPad than like android tablets that have broken like within a week. So, I would definitely recommend it. As far as quality goes, they are definitely worth the price. Plus, when you buy it you automatically get apple care and everything like that. And you can take it in to any apple store, and if anything is wrong with it they'll pay, they'll fix it for free, you just have to pay for any parts. And normally there's not any parts that break. Um, it's great for mail and everything like that. Yeah, and as far as facebook, any social media like that. And then games to kill time between class is kinda nice.

...

Other conversations are largely tangential to the purchasing decision. In some cases this is strategic; before the following conversation, for example, the seller had told us that her plan was to "chat with them and try to get them to connect with me so they will spend more:"

Seller: Oh yeah, sorry, I don't know about it, bad reception in the building or something. Yeah, so the whole thing of it is that, willing to see how much you would or how much you would be willing to spend to buy a basically a future which is a \$500 check for January 1st 2014.

Buyer: Yeah, yeah and I know a little bit about the time value of money, so.

Seller: No, I understand, so do I. Do you have a job?

Buyer: Yeah, I work part time for Ohio high school athletic association.

Seller: Oh, really thats cool. What do you do for them?

...

Notice how the seller quickly changes the topic of conversation from the decision at hand towards personal details in the hopes of establishing a better connection with the buyer.

Finally, some sellers were not above lying: in this example, the seller (successfully) misleads the buyer about a key detail of the experimental design:

Seller: So I was thinking we could find a price to make sure that we both win...
 Buyer: Mhmm, okay.
 Seller: So, what's your, basically, your bottom line, like, how much would you be willing to pay for this?
 Buyer: Well, how much does it, I guess, retail for?
 Seller: It retails for 500.
 Buyer: Okay, 500 dollars? Yeah. I'm thinking, probably then, my bottom would be [long pause] I mean I would be willing to pay 250 dollars for it.
 Seller: Well, yeah, but, [inaudible]
 Buyer: Ok, well, the highest I'd be willing to go, would probably be about 325.
 Seller: Ok, well, see there's another thing. If we sell it for, we can't make anything if we sell it for less than 360.
 ...

3.4 Researcher Predictions

Seven of the eight researchers whom we invited agreed to participate in our “prediction challenge.” We coded the 35 questions they proposed (Appendix B), which together provide a snapshot view into how social scientists currently understand persuasion. Two features are noteworthy.²²

First, congruence across researchers is low. To the human eye they appear to propose very different measures, with only three items appearing more than once: the number of words in the conversation (3×), the number of times the seller uses the buyer's name (2×), and how likeable the seller is (2×). Of course, variables could appear different but still capture similar statistical information. We therefore calculate how well a given variable proposed by one researcher is predicted by the variables proposed by a second researcher. This metric would be 100% if all researchers were in fact proposing the same underlying model, and 2% in expectation if their variables were completely unrelated (due to chance correlation). In our data the average value of this statistic is 26%, implying at most modest agreement. This suggests to us a relatively young field of research that has yet to converge on a consensus view of how persuasion works.

Second, researchers collectively expect both self-interest and other-regard to play a role in persuasion.²³ Of the thirty-five items, twenty one were coded by at least two of three research

²²We measure inter-coder reliability using Krippendorff's alpha, a statistic that is roughly one minus the average distance between coded values divided by the average distance between uniformly randomly coded values. A 1 indicates perfect reliability, and a 0 no better than randomly coded values. Many of our items proved difficult to code: Average alpha is 0.32 and median 0.26, with values ranging from -0.07 to 0.83. Since coder disagreements were not uncommon, to reduce noise and attenuation bias in further analysis, we collected more data when the coders disagreed. When the two coders disagreed on a subjective item - such as “how likeable do you find this pair to be” - a third coder from the same subject population also coded the item. Continuing to use subjects from the Ohio State subject pool for subjective items provided us the best estimate of how the conversational tactics were perceived in the conversation as it occurred. In such cases, we use the median value (which is “majority rule” for binary items) throughout the analysis. When the disagreement occurred on an objective item - such as the gender of the seller - then our research assistant provided her opinion, which we subsequently used. Of course, even with this second round of coding, any test of fit is always a joint test of the concept underlying the metric and the precision with which it is measured.

²³Among the researchers, we had both economists and psychologists. The economists' items did not differ meaningfully from the predictions of the psychologists in their focus on self-interest or other-regard, nor did

assistants as capturing appeals to other-regard, and fifteen items as capturing appeals to other-regard.²⁴ Researcher 3, for example, focused on self-interest, proposing a series of questions about product attributes (e.g. “How many statements of the form ‘you can use it for X’ does the seller make?”) and their interaction with a measure of trustworthiness. Several other researchers, meanwhile, asked questions about the likeability (or dis-likeability) of the seller, or about features like language style matching which are thought to promote liking.

4 Results

We now present the main empirical results. We conduct hypothesis tests non-parametrically using clustered rank-sum tests except where we need to condition on additional variables, in which case we use regression models based on the following specification:

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

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. Note that 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 ($p < 0.001$).

4.1 Does Communication Matter?

To motivate the rest of the analysis, we first examine whether communication affects buyer behavior at all. The answer is not obvious, since sellers have a material conflict of interest so that buyers might disregard everything they say as “cheap talk.” We find, however, that communication does alter valuations. Figure 1 shows the means of absolute changes in valuations with commission for the communication and no communication conditions. The mean absolute change in valuation under communication is \$49 (for all 15 sessions) significantly and substantially larger than the analogous \$24 without communication (clustered rank-sum $p < 0.01$).²⁵ On the extensive margin, 56% of buyers change their total valuation in the communication group and 31% in the no communication arm. Thus, while a sizeable minority of buyers are entirely unaffected by communication, the majority are, with some affected substantially. Further, though some of the variation in valuations seems to come from the effects of reflection, as measured in the No Communication arm, these changes fall short of explaining the changes we observe. The conversations are causing product valuations to change.

Result 1. *Communication affects buyer valuations.*

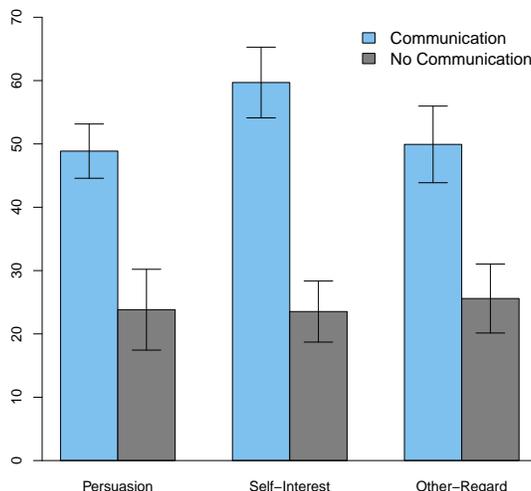
We also find that sellers benefit from communication. In principle it is possible that communication shifts some valuations up and others down without leading to better outcomes

either set outperform the other in predicting variance in outcomes for either pathway of persuasion (unreported).

²⁴The coding was not mutually exclusive, and two items were coded as intending to capture both pathways. No item was coded as intending to capture neither channel.

²⁵We are underpowered to perform the analysis within the last three sessions alone.

Figure 1: Absolute changes in valuation driven by communication



Mean of absolute valuation changes, by experimental condition.

for sellers on average – for example, if we think of communication as revealing a noisy signal that buyers use to update their beliefs, there is no reason to expect mean valuations to shift in either direction. What we find, however, is that valuations increase on average. The average change in buyers valuation with commissions is \$14 under communication (in the full 15-session sample), significantly different from the mean change in no communication of \$3 (clustered rank-sum $p = 0.01$) (See left-most bars in Figure 2). Also worth noting is though the No Communication treatment has significant absolute changes, we observe no systematic increase or decrease in mean valuations.

Result 2. *Sellers benefit from communication.*

One potential concern is that we observe valuation changes simply because buyers are confused or inattentive, and give inconsistent answers before and after the phone call. The facts that valuation changes are larger in the communication group than in the no communication group, and that valuations increase on average, are both direct evidence that this is not the sole explanation. We can also measure and control directly for subject confusion and inattention. We flag potentially confused or inattentive buyers in three different ways. First, we flag buyers who self-reported after conversations that they thought the instructions were unclear. Second, we had our coders flag buyers (and sellers) who they thought seemed confused about the rules of the game based on conversation transcripts. Third, we flag buyers who ever gave a non-monotonic response in the willingness to pay elicitation (e.g. willing to pay \$100, but not willing to pay \$80). Regardless of which definition we use, our results are qualitatively similar if we restrict the sample to subjects who were not confused or inattentive, and absolute

changes under communication are always significantly larger on average than those under no communication (results available on request).

4.2 Which Pathways Matter?

Result 3. *Persuasion works through both self-interest and other-regard.*

We turn next to our first core research question, whether other-regard matters for persuasion. We find evidence that both self-interest and other-regard matter. The second and third sets of bar graphs in Figure 1 illustrate this, plotting the means of absolute changes for the two components of persuasion. The middle bars show the mean absolute change in self-interest is \$60 compared to \$24 without communication (clustered rank-sum $p < 0.01$), and as in the right-most bars, the mean absolute change in other-regard is \$50 compared to \$26 in no communication (clustered rank-sum $p < 0.01$).²⁶ To quantify more precisely how much self-interest and other-regard contribute to variation in persuasion, we ask what proportion of the overall variation is eliminated by conditioning on each component.²⁷ We find that conditioning on self-interest reduces variance more than conditioning on other-regard, though 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.²⁸

For a less structured view of the data, and to check that averages are not unduly influenced by outliers, Figure 3 plots the full distribution of valuation changes. We show the change for each buyer-product pairing as a vector in (self-interest, other-regard) space, with the communication condition plotted at left and the no communication at right. We see an evident difference between the experimental conditions, with zero net change (vectors lying along $y = -x$) in 44% of conversations v.s. 69% of cases without conversations. We also see a general tendency towards higher valuations under communication, as there are more vectors lying northeast of $y = -x$ than southwest (36% v.s. 21%, compared to 13% and 18% without communication). Finally, both dimensions of persuasion are clearly in play, as changes are not concentrated along either axis. We count more increases than decreases in perceived self-interest (37% of vectors to the right of $x = 0$ versus 22% to the left) but see no systematic pattern for other-regard (25% of vectors above $y = 0$ and 28% below).

²⁶Though the seller’s surplus from a sale is large relative to the buyer’s, this result is not necessarily inconsistent with inequity aversion since the buyers expected gain from participation is larger than the sellers.

²⁷Specifically, we estimate

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

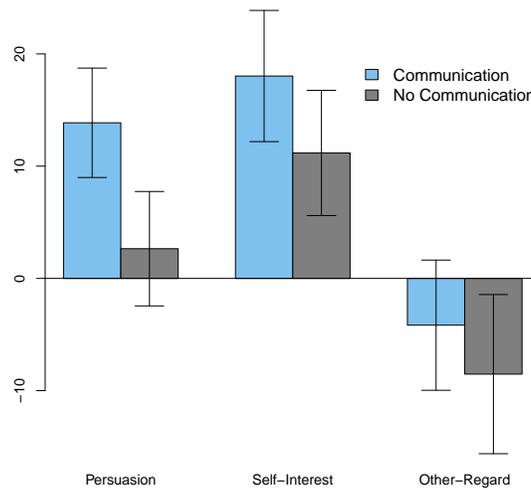
to measure the proportion of total variance attributable solely to self-interest, and

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

to measure that attributable solely to other-regard. Note that simpler variance decomposition procedures do not have a meaningful interpretation here given the non-independence of the two components.

²⁸Restricting the sample to subjects who self-report no confusion, or to pairs where both were coded as understanding the rules of the game, or to buyers who only gave monotonic responses does not alter these conditional variances substantially. In all cases, 15-19% of variance can be attributable to other-regard alone, and 28-36% can be explained by self-interest alone.

Figure 2: Communication increases valuations; changes in self-interest are mean positive



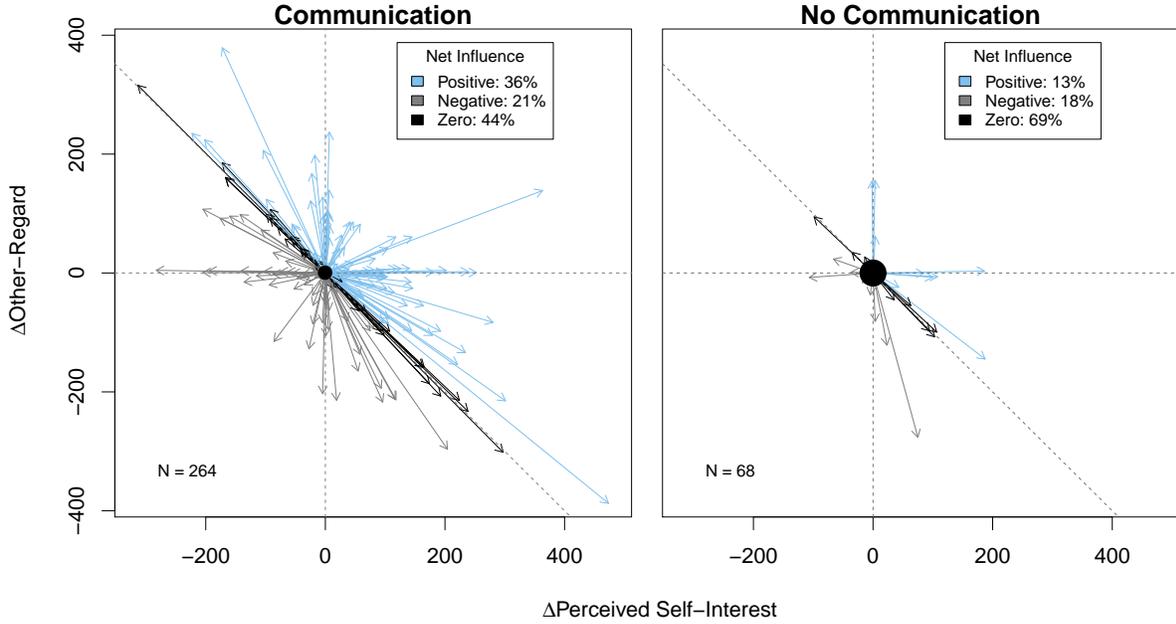
Plots the mean change in total persuasion, self-interest, and other-regard for both the no communication and communication conditions. Vertical bars depict standard errors allowing for clustering by buyer and by seller.

Indeed, changes in other-regard tend to cancel each other out on average. While the average change in perceived self-interest is \$18 and significantly different from zero, the average change in other-regard is -\$4 and insignificant (Figure 2). The same is true in the restricted sample: changes in self-interest net to a mean change of \$20, significantly different from zero (clustered t-test $p < 0.01$), comprising 115% of total persuasion, while the changes in other-regard roughly cancel each other out, netting a mean change of -\$2.64, insignificantly different than zero (clustered t-test $p = 0.7$).

At face value this might suggest that other-regard is simply “white noise” in the background, while self-interest is the main pathway for endogenous persuasion. For example, it could be that some buyers find some sellers more likeable than expected and others less so, but that sellers cannot or do not intentionally manipulate buyers regard for them. Indeed, many sellers explain their strategy solely in terms of buyer self-interest – they plan 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,” for example.

Yet many other sellers do report strategies that focus on other-regard. These sellers plan to “get the customer on my side,” to “build a rapport,” and to “get on buyer’s good side,” and so on. Some even mention specific other-regarding tactics from the literature: for example, “remember the person’s name! People love when others remember their names.” To analyze how such strategies fared, we had them coded by three undergraduate research assistants who independently assessed whether each referred to other-regard, self-interest, or contained any plan at all. Coders saw only strategy statements and subject identifier codes. We code a strategy as having a feature (e.g. targeting other-regard) if at least two of three coders identify it as such. Overall we code 19% of statements as targeting other-regard, 30% targeting self-interest, and 55% containing any plan. This coding appears consistent with our independent coding

Figure 3: Communication increases more valuations; net influence frequently driven by changes in self-interest



Plots the changes in self-interest and other-regarding valuation components as defined in Equation 1 for every buyer-seller pairing, for both the communication (left figure) and no communication (right figure) conditions. Each vector represents one buyer-seller pairing, with initial valuations normalized to the origin; vectors point to valuation changes, with jitter added. The x-axis denotes changes in perceived self-interest, and the y-axis denotes changes in other-regard. Vectors that lie northeast of the $y = -x$ line (light blue vectors) denote positive net influence; vectors lying east of the y-axis denote positive changes in perceived self-interest; vectors lying north of the x-axis denote positive changes in other-regard.

of the conversation transcripts themselves: sellers whom we code as targeting other-regard use the buyer’s first name directionally more often (*n.s.*), have more personal conversations ($p = 0.01$), and have more tangential conversation not focused on the product ($p = 0.03$).²⁹

We find that sellers who use other-regarding strategies perform significantly better than their peers. Table 2 summarizes the relationship between performance and strategy. The third column shows that sellers who use strategies that target other-regard succeed in 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, as the difference is small and insignificant. The first column shows the net result: 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.³⁰

²⁹ p -values from OLS regressions controlling for product fixed effects, clustering at seller level.

³⁰We also examine the strategy suggested by Dziuda (2011), that sellers should build credibility by saying something negative about the product. We find mixed evidence in support. Using RA coding of conversations, buyer valuations are significantly negatively correlated with negative seller statements, but significantly positively

Table 2: 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. “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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Result 4. *A minority of sellers plan to target other-regard and outperform those who do not.*

One question these results raise is whether sellers can endogenously increase other-regard. It could be that the sellers who report other-regarding strategies are really just intrinsically more likeable, for example, and that we measure the effects of their personality as opposed to their choices. While we cannot address this question conclusively without experimentally manipulating strategies, we can shed some light by examining sellers who change strategy across interactions. We do this in the last three columns of Table 2, where we add seller fixed effects to the model. The resulting estimates 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 because these specifications also include product fixed effects they are not simply capturing variation induced by the nature of the good being sold. It thus appears that other-regarding strategies are more successful than alternatives for the subset of sellers who use them.

4.3 How Does Persuasion Vary with Context?

We next examine how persuasion varies with the nature of the good being sold. Recall that we chose objects for sale to vary information asymmetry, with sellers better-informed than buyers about the tablet computers but symmetrically informed about the lottery and post-dated check (“abstract goods”). We therefore compare outcomes across these categories.

Consistent with our priors, sellers were more successful at persuading buyers to buy tablets. Figure 4 breaks down average valuation changes, overall and for the two pathways, for the tablets and the abstract goods. Valuations increase by \$25.1 for the tablets compared to \$9.6 for the abstract goods (rank-sum, clustered at buyer and seller level, $p = 0.03$). This is consistent with the view that information asymmetry creates greater scope for persuasion.

correlated with the interaction between positive and negative statements. We cannot replicate this pattern, however, using buyer self-reports on what they heard sellers say.

Yet the valuation increase for abstract goods is statistically significant: sellers benefit from persuasion even when selling lottery tickets and post-dated checks about which they have no special information. This result is hard to reconcile with models of persuasion in which buyer preferences are self-interested and fixed.³¹

Result 5. *Persuasion varies with the good sold, but is significant even for abstract goods.*

Perhaps persuasion is feasible even without asymmetric information because sellers can exploit other-regard? Consistent with this view, Figure 4 shows that other-regard plays a larger role for abstract goods than for tablets. The mean change in self-interest is significantly higher for tablets (\$34 versus \$6, rank-sum clustered at buyer and seller level, $p = 0.05$). In contrast, other-regard actually decreases on average for tablets, while increasing insignificantly for abstract goods (\$8.9 versus $-\$3.6$, rank-sum, clustered at buyer and seller level, $p = 0.22$). Broadly speaking, these results show that which pathways matter depend on the good being sold (or more generally, the decision being made). More specifically, they also suggest that self-interest plays a larger role when information is less symmetric. One caveat, however, is that we cannot reject the hypothesis that the same patterns hold in the no communication condition, due to imprecise estimates. Given this we consider the difference in pathways only suggestive.³²

We next examine how persuasion varies with the identities of the buyer and seller. We are particularly interested in what differentiates effective sellers; we know (above) that sellers who target other-regard outperform, but want a more general characterization. To examine this we exploit the panel structure of our data, which allows us to estimate the distribution of seller (and buyer) fixed effects and then ask (a) how much better the best sellers are than their peers, (b) on which pathways the better sellers distinguish themselves, and (c) what traits or behaviors correlate with being good at persuasion.

To our surprise, however, we find that seller effects are poor predictors of persuasion. In fact, buyer effects have much more explanatory power. Figure 5 depicts this, showing the adjusted R^2 s from a series of regressions, one per column, of persuasion ($WTPC_{bsp}^1 - WTPC_{bsp}^0$) on a particular battery of fixed effects.³³ Order and product effects are fairly unimportant for explaining the *change* in valuations. Seller effects are even less relevant, predicting less variation than would be expected purely by chance (adjusted $R^2 = -0.07$). They are all noticeably less predictive than buyer effects (adjusted $R^2 = 0.07$). For reference, only the prediction challenge winner was able to explain more variance than buyer fixed effects alone. Researchers as a group did not perform well, predicting a statistically insignificant 2% of the variation (F -test $p = 0.15$).

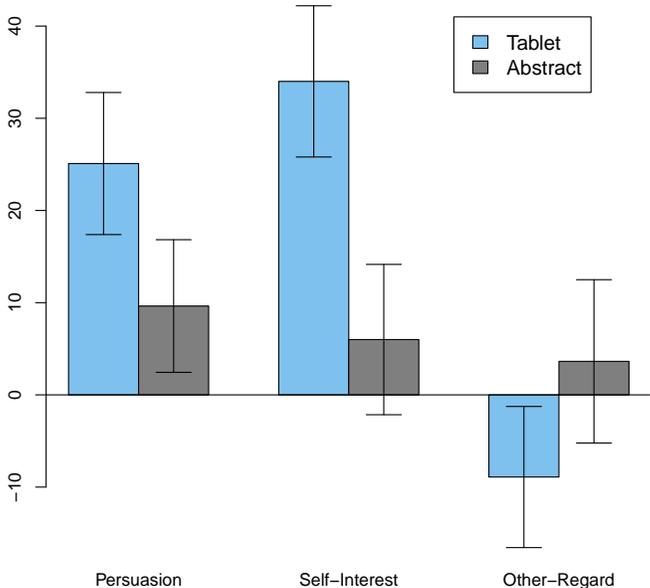
We draw similar conclusions when we conduct inference around these means. We test statistical significance by estimating models based on (2), augmenting with various sets of

³¹A more generic F -test for equality in mean persuasion across products is marginally insignificant ($p = 0.11$).

³²Conversations where the seller independently owns the tablet have directionally lowered valuations compared to other pairings with a tablet as the object. This could be because product ownership is not important, or because most sellers own an iPad, and valuations were more fixed for this tablet compared to the Surface.

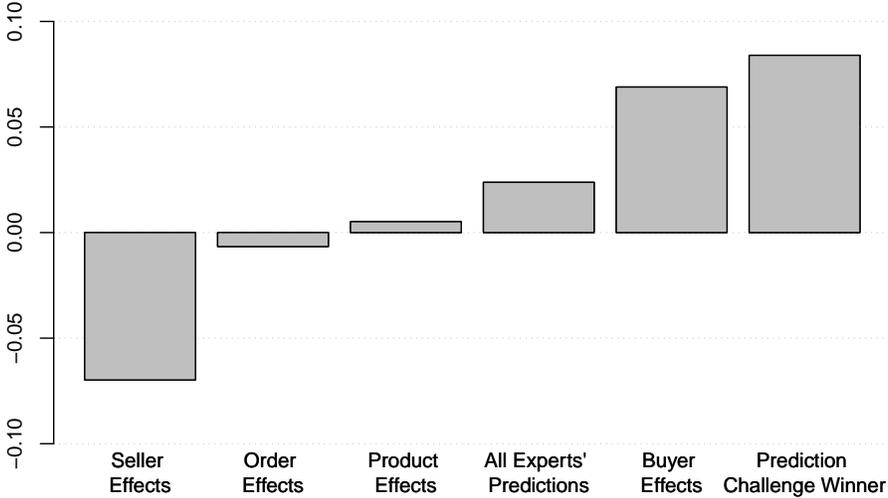
³³Recall that the “adjustment” in adjusted R^2 lets us compare goodness of fit across models with different numbers of predictors by benchmarking explanatory power against that one would expect from an equal number of randomly generated variables.

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.

Figure 5: 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. Researcher predictions for “All Predictions” and “Prediction Challenge Winner” interacted with product fixed effects.

fixed effects, and in each case conducting an F -test of the restriction that the included fixed effects are jointly zero, using both standard asymptotics and randomization inference.³⁴ We cannot reject the null that seller effects are jointly zero ($p = 0.87$), but easily reject the null that buyer effects are zero ($p < 0.01$). Figure A-2 plots the distributions of (absolute) effects being tested. Each point represents mean absolute change in valuation for one individual, ordered by mean change in valuation. The distribution of seller effects is flat relative to that for buyers, whose distribution increases more rapidly in both directions away from the nadir.³⁵

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

Of course, we know from above that sellers are not all the same – those who target other-regard outperform those who do not. That test is better-powered than an omnibus test of seller fixed effects as it tests against a specific alternative hypothesis. Rather, the result here says that buyer identities are generally more predictive of persuasion than seller identities, or more generally that the evidence for *persuadable* types is stronger than the evidence for *persuasive* types.

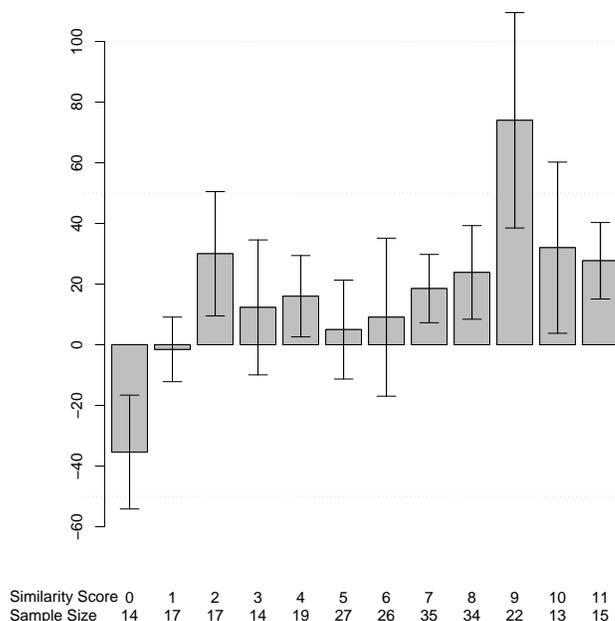
We also explore whether the *match* between buyer and seller matters, conditional on their individual characteristics. We measure homophily, or buyer-seller similarity, in two complementary ways. First, we use buyers’ self-reports of their perceived similarity with the seller. The benefit of this measure is that it is comprehensive; the drawback is that it is subjective, and buyers might feel a need to answer in a way that rationalized their valuation responses. We therefore complement it by also testing directly for effects of a match on the observable characteristics of gender and nationality. These are incomplete measures of homophily, but objective. Overall both measures appear sensible. Perceived similarity is distributed fairly uniformly 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. 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 $p < 0.01$).

We find that homophily strongly predicts persuasion. Figure 6 illustrates this for self-reported similarity, while Table 3 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

³⁴Standard asymptotic inference for this statistic requires that observations per participant approach infinity, which may not be tenable in our short panel. Under the null hypothesis that changes in valuation are independent of the individuals involved, however, the distribution of our data is unaffected by permuting individual identifiers. We therefore conduct randomization inference by drawing 1,000 permutations of individual identifiers, re-estimating the F -statistic after each, and using the resulting distribution to estimate the sampling distribution of our original F -statistic (Lehmann and Romano, 2005). The p -values we obtain in this manner are essentially identical to those we obtain using standard asymptotics.

³⁵The data do not make it clear whether the buyers differ in propensity to be persuaded to like a product or a person. F -tests for buyer effects on changes in both self-interest ($p = 0.19$) and other-regard ($p = 0.53$) are insignificant. The corresponding figures for sellers are $p = 0.64$ and $p = 0.38$.

Figure 6: 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.

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 ($p = 0.04$).³⁶ Only for nationality is this pattern weaker; persuasion is higher by \$45.4 for same-nationality pairs but this is not statistically significant.³⁷

Result 7. *Persuasion increases in buyer-seller homophily, and in particular gender match.*

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.^{38,39}

One open question about these homophily results is whether they reflect the exogenous

³⁶We estimate similar effects of male-male and female-female matches.

³⁷While two non-US citizens are not necessarily well described as “matched,” anecdotally the probability is high that both are East Asian, and in any case such pairings are too rare (4%) to be driving our results.

³⁸We also find that similarity continues to predict persuasion even after we condition on a measure of trust coded from the transcripts (not reported).

³⁹Concurrent with our work, He et al. (2015) also find that gender-matched pairs cooperate more in social dilemma games that involve communication.

Table 3: Buyer-seller homophily predicts persuasion through both pathways

	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	1.13	19.79
Seller Non-American							(9.46)* -11.61	(16.33) -30.16	(19.97) 16.91
N	210	210	210	220	220	220	(7.42) 216	(21.93) 216	(15.16) 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. "Similarity score" is the buyer's perceived similarity with the seller on a 1-11 scale. Standard errors clustered by buyer and by seller in parenthesis. Statistical significance is denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

impacts of buyer and seller types, or whether buyer perceptions of homophily are also shaped by endogenous seller behaviors. For example, at least one seller had an intentional strategy of identifying things he had in common with the buyer. Can such strategies work?

We examine this in two steps. We first identify conversational behaviors that predict buyers' perceptions of homophily. Of thirty-four items we coded, six significantly predict homophily (Table 4, Columns 1 & 2).⁴⁰ They are items that seem intuitively likely to be associated with similarity – for example, frequently using the buyer's name, getting personal, and so on. This pattern seems consistent with the idea that conversation can endogenously shape perceptions of homophily. Since conversations might have these features precisely when the participants have similar exogenous traits, however, we examine whether perceived similarity in a given interaction is predicted by a seller's (or buyer's) behavior in their *other* three conversations. We find that, for all six items, mean seller behavior in other conversations does predict homophily in the conversation at hand (Table 4). Notably, this does not hold for buyers, suggesting that it is seller behavior that drives perceived homophily.

Result 8. *Seller behavior can increase buyer-seller homophily.*

One further concern with this interpretation is that some sellers could be both more likely to have conversations with these features and to have traits that lead buyers to perceive them as similar. For example, if men were over half the subject pool and also more likely to refer to each other by name, we might erroneously conclude that the use of first names increases perceptions of homophily when in fact gender does so. We cannot control for all such potentially confounding traits, of course, but do show in Columns 2, 4, and 6 of Table 4 that controlling for the major covariates we do observe (gender and nationality) has little impact on the results. While not fully dispositive, this strongly suggests that it is endogenous tactics and not exogenous traits that are driving the result.

4.4 How do Self-Interest and Other-Regard Interact?

Finally, we investigate how the two pathways of persuasion interact with each other. In principle there could be either mechanical or causal interactions. Mechanical interactions would arise if sellers face constraints in terms of time, buyer attention, energy, etc. which force them to choose between influencing self-interest or other-regard. For example, time spent bonding over last night's football game cannot be spent discussing the product. Causal interactions could arise if a given seller behavior affects both self-interest and other-regard. These effects could be complementary or offsetting. For example, spending time describing positive features of a product could come across as helpful and engender feelings of reciprocity in the buyer, or as pushy and engender resentment. Establishing things that the parties have in common could affect both other-regard and the credibility of advice about the buyer's self-interest. In fact, the only interaction we could not plausibly rationalize a priori is how building other-regard could detract from persuasion through self-interest.

We first examine the raw correlation between changes in self-interest and other-regard. We

⁴⁰We control here for total number of words, decreasing the total number of coded items available from thirty-five to thirty-four.

Table 4: Conversation tactics affect perceived homophily

<i>Value of indept. variable:</i>	Dependent Variable = Buyer's self-reported similarity with seller					
	From same conversation as reported similarity		Avg from seller's other three conversations		Avg from buyer's other three conversations	
Times seller uses buyer's first name	0.61 (0.29)**	0.52 (0.26)*	0.76 (0.30)**	0.67 (0.26)**	-	-
How personal does buyer get	0.60 (0.25)**	0.46 (0.25)*	0.85 (0.21)***	0.60 (0.23)***	-0.19 (0.43)	-0.20 (0.43)
How warm or likeable is pair	0.89 (0.39)**	0.71 (0.38)*	1.10 (0.48)**	0.86 (0.45)*	0.47 (0.63)	0.50 (0.70)
Fraction words spoken by buyer	2.65 (1.58)*	2.91 (1.74)*	3.28 (1.87)*	2.97 (1.93)	-1.28 (3.02)	-0.53 (2.94)
Number overlaps/interruptions	0.07 (0.02)***	0.06 (0.02)***	0.07 (0.02)***	0.06 (0.02)***	0.03 (0.04)	0.04 (0.04)
First name, 'you', or 'your'	0.89 (0.25)***	0.77 (0.24)***	0.95 (0.33)***	0.79 (0.29)***	-	-
Gender-match control?	No	Yes	No	Yes	No	Yes
Citizenship-match control	No	Yes	No	Yes	No	Yes

Each cell is a separate OLS regression. We regress buyer's post-conversation, self-reported similarity with the seller after the conversation on one variable from transcription codings (noted in the far left column), and number of words in the conversation. Columns 1 and 2 report coefficients from regressing reported similarity on the transcription variable value from the same conversation; Columns 3 and 4 report coefficients from regressions using the mean value of that variable for the seller's *other* three conversations; Columns 5 and 6 report coefficients from regressions using the mean value of the item for the *buyer's* other three conversations. We include every coded transcript item (i) that significantly predicts reported similarity while controlling for number of words spoken and (ii) that is positively, significantly correlated with the mean value of that same item for the other three conversations (for the seller for columns 1-4 and for the buyer for the final two columns). Standard errors clustered by buyer and by seller in parenthesis. Statistical significance is denoted as: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

find a strong, significant negative correlation $\rho = -0.64$. This could be biased by measurement error, however, since the change in self-interest also enters negatively into our calculation of the change in other-regard.⁴¹ We therefore project changes in self-interest and other-regard onto a full set of conversation features and product fixed effects and examine the correlation between the fitted values. This limits the analysis to the component of persuasion not orthogonal to these regressors, but also eliminates any spurious negative correlation due to classical measurement error.⁴² We find that predicted changes in self-interest and other-regard are in fact even more negatively correlated, $\rho = -0.70$, and this holds for all four products individually (iPad $\rho = -0.79$; Surface $\rho = -0.79$; Check $\rho = -0.68$; Lottery $\rho = -0.52$). This suggests a sharp trade-off between the two pathways.

This pattern does not appear to be driven by mechanical substitution between tactics in short conversations. First, none of the conversations had to be stopped because of the ten-

⁴¹Recall that increases in other-regard are measured by changes in total valuations less changes in valuations for the good (Equation 1). The covariance of the pathways is thus $\text{Cov}(\Delta WTPF_{bsp}, \Delta WTPC_{bsp} - \Delta WTPF_{bsp})$, equivalent to $\text{Cov}(\Delta WTPF_{bsp}, \Delta WTPC_{bsp}) - \text{Var}(\Delta WTPF_{bsp})$. Hence if $WTPF_{bsp}$ is observed with error ($\text{Var}(\Delta WTPF_{bsp})$ is inflated), this term is biased downwards.

⁴²Formal derivation available on request.

minute time constraint. Second, we obtain similar results regardless of conversation length, even in the longest conversations in which one might expect the mechanical constraints to be less binding. If anything we see a sharper trade-off in longer conversations: by quartile of conversation length, the correlations are $\rho = -0.54$, $\rho = -0.76$, $\rho = -0.73$, and $\rho = -0.70$ respectively. It thus seems unlikely that sellers simply run out of time, and more likely that their behaviors have offsetting effects. The interaction of the pathways indicates a theory-driven approach will require more than the summation of the two forces. Further, the result provides a challenge for future work of determining why the give-and-take exists. It could be reflect real constraints – for example, seller tactics may work in their intended pathway but backfire in the other. It could also work through mental accounting, as is suggested for example by Hastings and Shapiro (2013) study of consumer responses to changes in gasoline prices.

5 Conclusion

In this paper we define and apply an experimental framework for studying the pathways of persuasion: in particular, whether and when persuasion works through appeals to self-interest 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. Variation in persuasion is driven less by who is selling than by who is buying, and his also strongly influenced by homophily. Finally, sellers face a trade-off, with gains on one pathway generally offset in part by losses on the other.

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. This seems especially interesting given that we do *not* find significant variation in seller performance in our sample of novices.

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.

Future work could also shed more light on what drives the trade-off between pathways. Understanding this interaction could be crucial for theoretical work introducing other-regard into persuasion. While we cannot conclusively answer this question, it seems likely that selling the good comes at the expense of being liked, rather than the other way around. First, we have no reason to hypothesize a priori that building other-regard could causally decrease a buyer's perceived self-interest in the product. Second, sellers who successfully focus on increasing other-regard did not generate decreases in self-interest (Table 2). This suggests that the negative correlation between pathways we observe may reflect buyer dislike for sellers who try too aggressively to manipulate self-interest. Subsequent work could make this point conclusive as well as isolate precisely what about appealing to self-interest the listener finds so unlikeable.

Finally, we conjecture that our results offer a new perspective on bargaining. In real-world bargaining over the sale of a car (say), the participants can both negotiate a price and also seek to persuade each other of the car's value. Experimental studies of bargaining typically fix value, while letting participants negotiate price; our experiment fixes price (stochastically), while letting participants discuss value. It would be interesting to study the role of self-interest and other-regard in a bargaining environment where both factors were in play.

References

- Andreoni, James and A. Abigail Payne**, “Charitable Giving,” in Alan Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *Handbook of Public Economics, Volume 5*, 2013.
- **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.
- Antioch, Gerry**, “Persuasion is now 30 Percent of US GDP,” *Economic Roundup: The Treasury of the Australian Government*, June 2013, (1).
- Becker, Gary S and Kevin M Murphy**, “A Simple Theory of Advertising as a Good or Bad,” *The Quarterly Journal of Economics*, November 1993, *108* (4), 941–64.
- Bellemare, Charles and Bruce Shearer**, “Gift giving and worker productivity: Evidence from a firm-level experiment,” *Games and Economic Behavior*, 2009, *67*, 233–244.
- Bohnet, Iris and Bruno S Frey**, “The Sound of Silence in Prisoner’s Dilemma and Dictator Games,” *Journal of Economic Behavior and Organization*, 01 1999, *38* (1), 43–57.
- Bolton, Gary E and Axel Ockenfels**, “ERC: A theory of equity, reciprocity, and competition,” *American Economic Review*, 2000, pp. 166–193.
- Brandts, Jordi, Gary Charness, and Matthew Ellman**, “Let’s Talk: How Communication Affects Contract Design,” *Journal of the European Economic Association*, forthcoming.
- Buchan, Nancy R, Eric J Johnson, and Rachel TA Croson**, “Let’s get personal: An international examination of the influence of communication, culture and social distance on other regarding preferences,” *Journal of Economic Behavior and Organization*, 2006, *60*, 373–398.
- Burger, Jerry M, Nicole Messian, Shebani Patel, Alicia del Prado, and Carmen Anderson**, “What a Coincidence! The Effects of Incidental Similarity on Compliance,” *Personality and Social Psychology Bulletin*, 2004, *30* (1), 35–43.
- Byrne, Donn Erwin**, *The Attraction Paradigm*, Vol. 11, Academic Press, 1971.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez**, “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction,” *American Economic Review*, 2012, *102* (6), 2981–3003.
- 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 Matthew Rabin**, “Understanding Social Preferences With Simple Tests,” *The Quarterly Journal of Economics*, August 2002, *117* (3), 817–869.
- **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.

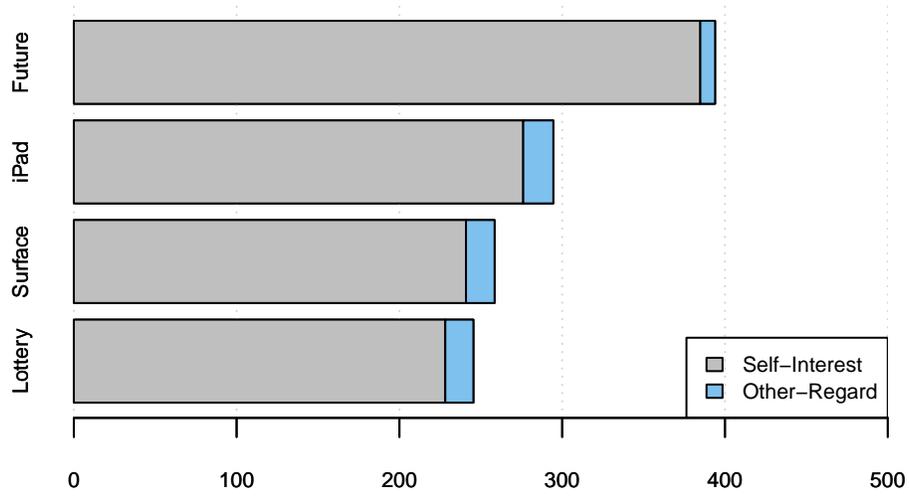
- , **Sanford Braver**, and **Stephen Lewis**, “Attributional Bias and the Easily Persuaded Other,” *Journal of Personality and Social Psychology*, 1974, *30* (5), 631–637.
- Cooper, David** and **John Kagel**, “Other-Regarding Preferences: A Selective Survey of Experimental Results,” in Alvin Roth and John Kagel, eds., *Handbook of Experimental Economics, Volume 2*, in press.
- 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.
- Dawes, Robyn M**, **Jeanne McTavish**, and **Harriet Shaklee**, “Behavior, communication, and assumptions about other people’s behavior in a commons dilemma situation.,” *Journal of Personality and Social Psychology*, 1977, *35* (1), 1.
- DellaVigna, 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.
- Emswiler, Tim**, **Kay Deaux**, and **Jerry E Willits**, “Similarity, Sex, and Requests for Small Favors,” *Journal of Applied Social Psychology*, 09 1971, *1* (3), 284–291.
- Fehr, Ernst** and **Klaus M Schmidt**, “A theory of fairness, competition, and cooperation,” *Quarterly journal of Economics*, 1999, pp. 817–868.
- 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.
- Hastings, Justine S** and **Jesse M Shapiro**, “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks,” *The Quarterly Journal of Economics*, 2013, *128* (4), 1449–1498.
- He, Simin**, **Theo Offerman**, and **Jeroen van de Ven**, “The Sources of the Communication Gap,” *SSRN Working Paper No. 2598441*, 2015.
- 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.
- Jiang, Lan, Joandrea Hoegg, Darren W Dahl, and Amitava Chattopadhyay**, “The Persuasive Role of Incidental Similarity on Attitudes and Purchase Intentions in a Sales Context,” *Journal of Consumer Research*, 02 2010, 36 (5), 778–791.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, October 2011, 101 (6), 2590–2615.
- Krueger, Alan B and Alexandre Mas**, “Strikes, scabs, and tread separations: Labor strife and the production of defective bridgestone/firestone tires,” *Journal of Political Economy*, 2004, 112 (2), 253–289.
- 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.
- Malmendier, Ulrike and Klaus Schmidt**, “You Owe Me,” *NBER Working Paper 18543*, 2012.
- Marmaros, David and Bruce Sacerdote**, “How Do Friendships Form?,” *The Quarterly Journal of Economics*, 02 2006, 121 (1), 79–119.
- McCloskey, Donald and Arjo Klamer**, “One Quarter of GDP is Persuasion,” *American Economic Review, Papers and Proceedings*, May 1995, 85 (2), 191–195.
- Miller, Dale T, Julie S Downs, and Deborah A Prentice**, “Minimal Conditions for the Creation of a Unit Relationship: The Social Bond between Birthdaymates,” *European Journal of Social Psychology*, 1998, 28 (3), 475–481.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer**, “Coarse Thinking and Persuasion,” *The Quarterly Journal of Economics*, 05 2008, 123 (2), 577–619.
- Murphy, Kevin M and Andrei Shleifer**, “Persuasion in Politics,” *American Economic Review, Papers & Proceedings*, May 2004, 94 (2), 435–439.
- 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.

- Rabin, Matthew**, “Incorporating fairness into game theory and economics,” *The American economic review*, 1993, pp. 1281–1302.
- Roth, Alvin E.**, “Bargaining Experiments,” in Alvin Roth and John Kagel, eds., *Handbook of Experimental Economics*, 1995.
- 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.

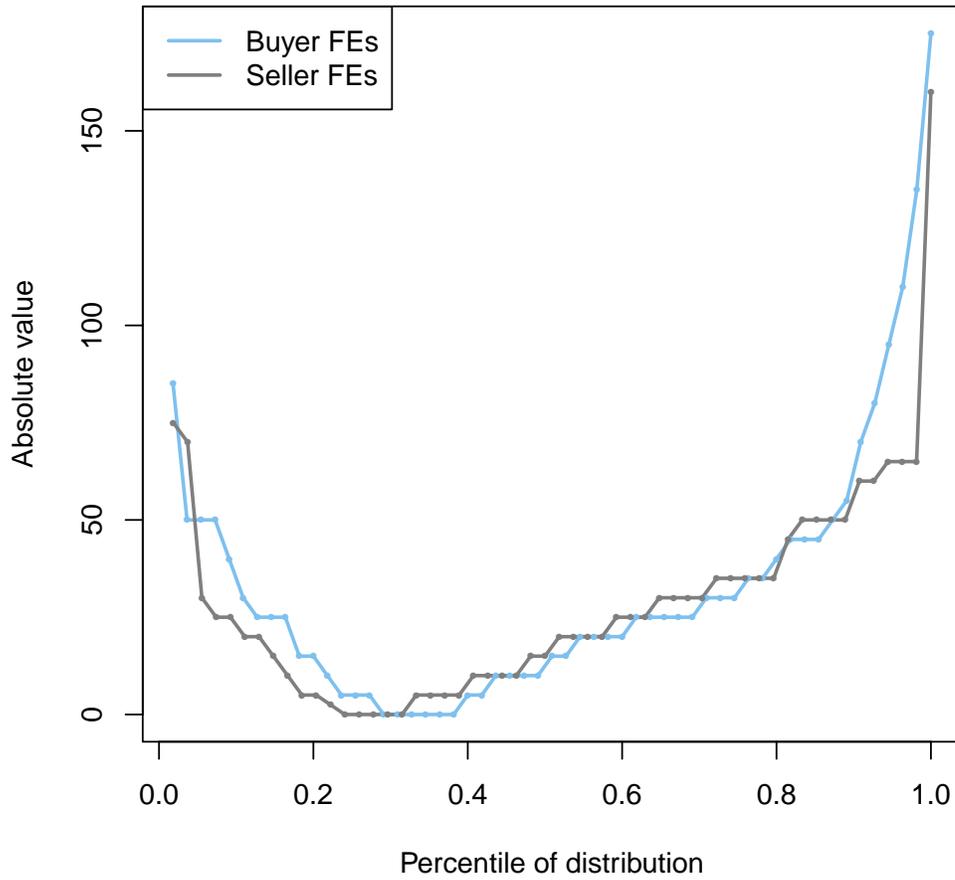
A Data Appendix

Figure A-1: Initial valuations show moderate 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").

Figure A-2: Distribution of participant fixed effects



Plots each study participants' absolute mean change in persuasion across all four conversations, sorted from smallest to largest before absolute value taken, for each seller (gray) and each buyer (blue), from all sessions without the no communication condition (1-11).

Table A-1: Participation by Session

Session	Buyers	Sellers	Conversations	No Communication
1	4	4	16	0
2	6	6	24	0
3	5	5	20	0
4	5	5	20	0
5	7	7	28	0
6	4	4	16	0
7	5	5	20	0
8	5	5	20	0
9	5	5	20	0
10	5	5	20	0
11	4	4	16	0
12	10	6	20	20
13	7	3	8	20
14	6	2	4	20
15	5	4	12	8
Totals	83	70	264	68

Table A-2: Participant characteristics

Variable	Buyers		Sellers	
	Mean	SD	Mean	SD
Age	21.24	1.86	21.21	1.80
Gender = Male	0.55	0.50	0.56	0.50
Weekly Expenditure	67.34	42.74	86.95	82.93
Voted	0.56	0.50	0.70	0.46
Owns an iPad	0.00	0.00	0.80	0.40
Used an iPad	0.88	0.32	0.94	0.24
Owns a Windows Tablet	0.00	0.00	0.24	0.43
Used a Windows Tablet	0.18	0.39	0.56	0.50
Owns an iPhone	0.54	0.52	0.68	0.47
Owns an iPod	0.76	0.46	0.80	0.40
Owns an Android Tablet	0.06	0.29	0.14	0.34

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 tablet.

B Researcher Predictions for Correlates of Persuasion

Here, we present the predictions from the persuasion researchers 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.

Researcher 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?

Researcher 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).

Researcher 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 $(4) \times [(1) + (2) + (3)]$

Researcher 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?

Researcher 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)

Researcher 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?

Researcher 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)