

## 1 Introduction

In *An Economic Theory of Democracy*, Anthony Downs (1957) proposes a solution to the problem of electorates who lack the interest, ability, or time to engage with politics. Essentially, he calls for a division of labor in which those who have the interest, ability, and time to engage with politics pay attention, and the rest of the public asks those people how they should vote and what their opinions on key issues should be. To avoid being misled, those who need this cheap information should seek out informants who have similar political preferences. Otherwise, the informant may advise them to support a candidate or policy that promotes the *informant's* preferences while undermining the recipient's goals. This sender-receiver framework is a logical solution to the problem of poorly informed electorates that has spawned a long research agenda – including many experiments, some of which we conducted (e.g., Krupnikov et al., 2020; Pietryka, 2016; Ryan, 2011b).

But this framework is an odd description of human interactions. The oddity derives from the motivations of the actors in the model. First, the theoretical communication is instrumental, initiated by the person who lacks information. If this person does not seek advice, the interaction would never take place (at least as described by Downs [1957]). Yet people who lack political information often *avoid* – rather than seek – political discussions (Carlson & Settle, 2022), suggesting that this framework mischaracterizes the motivations driving such conversations.

Instead, these conversations are more likely to be initiated by informants, since they are the ones who are most interested in politics (Huckfeldt & Mendez, 2008). In fact, informed people – rather than the uninformed who, according to Downs, *should* be initiating communication – initiate political communication all the time. Krupnikov and Ryan (2022) demonstrate that those who are “deeply involved” in politics are those who are most likely to talk – and post on social media – about politics. They hold more extreme political preferences and tend to be more affectively polarized (see also Klar, Krupnikov, & Ryan, 2018).

This framework also assumes these conversations are motivated by political concerns – that is, it assumes people seek informants who are best able to provide them with political information. Individuals *could* select these informants from their broader social networks (see, e.g., Eveland & Kleinman, 2013), but this assumption conflicts with most observed political discussion networks. Instead of conducting a motivated search for political informants, people more commonly encounter political discussion as a byproduct of everyday conversations (Minozzi et al., 2020). As a result, the political informants people rely on

most heavily are the same people they most frequently discuss other important matters with – often their close friends and family (Klofstad, McClurg, & Rolfe, 2009).

Further, in the basic Downsian model, the only way people can trust an informant is when they have shared political preferences. This is primarily because it is assumed that the informant’s main goal is persuading others to reach the same conclusions they have reached when two people with conflicting preferences should reach different conclusions. Certainly, individuals send biased messages about politics – this tendency is reflected in the importance of discussion network partisanship in voting decisions (Ryan, 2010; Sokhey & McClurg, 2012) – but it is unlikely that the Downsian model characterizes most political discussion: if it did, preferences would merge enough that disagreement in networks could not survive (Huckfeldt, Mendez, & Osborn, 2004). Instead, people often discuss politics for reasons *other* than information seeking or persuasion (Lyons & Sokhey, 2014): people often discuss politics to maintain social bonds, exchange perspectives, or just pass time (Carlson & Settle, 2022; Conover, Searing, & Crewe, 2002; Eveland, Morey, & Hutchens, 2011). Yet, in Downs’s framework, informants gain no satisfaction from such social motivations.

In sum, Downs envisions a world in which political discussion occurs because individuals who lack sufficient political information seek more-expert informants. And these informants care more about furthering their political goals than helping others, sharing their expertise, or maintaining healthy relationships. These exchanges do not seem to offer a “solution” to the problem of electorates who lack the interest, ability, or time to engage with politics. Nor do they seem to reflect what discussion actually looks like, where communication is often intended to fulfill motivations *beyond* political information seeking and is associated with a general interest in politics on the part of the person doing the talking.

And yet the experimental research on interpersonal communication supports Downs (1957). Participants tend to prefer more expert sources (Ahn & Ryan, 2015; Pietryka, 2016). They send biased summaries of the news (Carlson, 2019). They make worse decisions when receiving information from people with different preferences (Krupnikov et al., 2020; Lupia & McCubbins, 1998; Ryan, 2011b). In all, much of the experimental literature suggests interpersonal communication is indeed driven by partisan motivations (Ahn, Huckfeldt, & Ryan, 2014) with little aggregate enlightenment as a result (Jackman & Sniderman, 2006).

One critique of this experimental literature – and a potential explanation for findings that support Downs (1957) despite his odd assumptions – is that the

*Examining Motivations in Interpersonal Communication* 3

results in these previous experiments are only externally valid to contexts in which participants have motivations similar to the ones that Downs (1957) was talking about. A related critique is that the experimental designs either explicitly or implicitly motivate participants to think like Downs would expect them to. The experiments are often explicitly built around a model in which communication arises from information seeking rather than to fulfill other motivations. The researchers do not tell participants that trust only comes from shared political preferences, but since the instructions explicitly frame the experiments in political terms and emphasize preferences, they imply that preferences are a key element. It thus should be no surprise that participants' behavior is driven by these preferences (see Groenendyk & Krupnikov, 2021).

If we are going to build a model that extends to more common interpersonal discussion contexts than the one put forth by Downs and subsequent experiments, we need to do so explicitly. Thus, in this Element, we primarily discuss experiments in the behavioral economics tradition, incentivizing participants to consider motivations besides the ones typical in the Downsian perspective. After explaining our experimental framework in Section 2, we move beyond the information-seeking model of interpersonal communication (Section 3). There are incentives for participants to seek information, but also to *send* information. Further, we include treatments in which the motivation to acquire information from news sources is correlated with the desire to send messages – as Krupnikov and Ryan (2022) would suggest is fairly common. In Section 4, we introduce the possibility that individuals might have prosocial motivations in addition to their partisan and accuracy ones. That is, we create incentives to help others in the experiment, just as real-world informants are often motivated to help their friends and family.

The results suggest that motivations matter. When participants have incentives that align with Downs's model, interpersonal communication promotes the welfare of experts at the expense of the less informed. For instance, when the most interested also tend to have the most extreme preferences, moderates tend to be poorly represented (Section 3). On the other hand, the introduction of prosocial motivations in Section 4 shows greater potential for individuals to improve their political decisions via communication. Together, the results suggest that interpersonal communication is most likely to be effective as a political information shortcut when the motivations behind the communication are not political – especially when individuals receive less of a benefit from persuading others. This finding, especially when combined with the accumulated literature (see Krupnikov & Ryan, 2022; Minozzi et al., 2020; Settle, 2018), gives greater hope to the value of interpersonal discussion for collective civic capacity than discussion via social media.

The designs of our experiments are built off the basic design in other interpersonal communication experiments in the behavioral economics tradition (e.g., Ahn et al., 2014; Krupnikov et al., 2020; Pietryka, 2016) and are expanded upon in more detail in Section 2. In these studies, participants are brought into a laboratory in groups to participate in a mock election. Researchers can vary features of the election, such as the candidates' policy positions, the participants' policy positions, and how participants learn about these positions. Importantly, researchers provide incentives to the participants that determine what motivates participant behavior – whether the researchers know this is what they are doing or not (see Groenendyk & Krupnikov, 2021).

We provide an outline of the experiments in Figure 1 – we build on and edit this figure in our experimental descriptions in Sections 3 and 4. At the first stage in each election, the experimental factors are randomly assigned, as are the positions of two computer-generated candidates and the positions of each participant. In our experiments, these factors are designed to incentivize specific motivations – we discuss this after explaining how participants learn about the candidates. In each election, participants are voters who are each assigned an integer position on a numeric scale with no defined policy content. The participants must choose between two “candidates” who are not fellow participants but whose numbers represent the candidates' positions on the aforementioned scale. Participants gain money when the candidate closest to their position wins

**Stage 1. The election begins**

- Experimental factors assigned
- Candidates' positions assigned
- Participants' positions assigned

**Stage 2. Private information and first judgment**

- Participants view private information
- Each signal drawn randomly from uniform distribution centered on candidate's true position  $\pm 3$
- After viewing information, participants enter estimates of each candidate's position

**Stage 3. Social information and second judgment**

- Participants choose with whom to communicate
- Participants send information
- Participants enter updated estimates of the candidates' positions

**Stage 4. Vote**

- Participants vote

**Stage 5. Outcome**

- Participants learn winner of election and the rewards they received
- The winner of the election determined by majority rule

**A new election begins at Stage 1**

**Figure 1** An overview of the stages in an election experiment

*Examining Motivations in Interpersonal Communication* 5

the election. The candidates' positions are unknown to participants, but they can learn about these positions in two ways: private information and social information.

In Stage 2 of each election, participants can obtain **private information** that is noisy but unbiased on average, mimicking the acquisition of impersonal sources like news media. Like in the real world, participants' expertise varies – in the experiments, expertise is measured as the amount of private information the participant has received. After obtaining the private information, in Stage 3, participants can also obtain **social information** from fellow participants, mimicking interpersonal communication. Regardless of whether the information is private or social, each piece of information contains two integer signals indicating estimates of each candidate's position.

Subjects make three judgments. The first is an estimate of the candidates' positions after receiving private information (Stage 2). The second is a (potentially) updated estimate of the candidates' positions after receiving the social information (Stage 3). These two judgments allow us to see how interpersonal communication affects participants' beliefs about the candidates. The final judgment comes in Stage 4 when participants vote for a candidate. This is the most consequential judgment since it determines the payoffs of the election – though the exact incentives vary in each experiment, participants always receive more money when the victor is the closest candidate to their position. This winner is revealed in Stage 5, after which a new election begins at Stage 1.

Our key experimental factors are designed to manipulate motivations by randomly assigning payoffs that are not exclusively based on the election outcome. In the experiments featured in Section 3, subjects receive (or lose) money based on whether they participate in the interpersonal communication stage. The actual discussion payoff is randomly assigned at the individual level. At the group level, we randomly assign the correlations between the participants' discussion payoffs, positions, and expertise.

In Section 4, we incentivize prosocial considerations by paying some participants based on the decision of their discussion partners – that is, the senders are paid if the receivers “vote correctly” (Lau & Redlawsk, 1997). On one hand, this incentivizes senders to truthfully reveal to receivers how they believe the receivers should vote. On the other hand, if the senders and receivers have different preferred candidates, then the senders lower their probability of receiving the election payoff if they earn the prosocial payoff. Hence, it is not a forgone conclusion that subjects will be helpful in the prosocial treatment.

We obviously value the findings in previous experimental studies, but the new experiments discussed in this Element are important additions to the literature since nonpartisan motivations affect every stage of interpersonal political communication, from the decision to join a conversation to how people react after a conversation has ended. This point is a core contribution of Carlson and Settle (2022). We hope their work shifts attention to the importance of motivations, but more work is needed: since their impressive designs do not randomly assign motivations, they lack the ability to identify the effects of motivations or distinguish between the effects of different motivations (Leeper & Slothuus, 2014).

This is why we advocate for studying interpersonal communication motivations in laboratory settings as we have done here. Of course, the experiments described here are abstract, which limits their external validity. Further, the motivations we examine are not an exhaustive list of all motivations that may be relevant for interpersonal political communication – nor do they include all the motivations that can be experimentally manipulated. We therefore conclude this Element with a discussion of both the promise and perils of studying motivations in interpersonal communication via laboratory experiments.

## 2 Our Experimental Framework

Social influence, including interpersonal communication, is not an understudied aspect of society and there has been wide variation in approaches to its study. Some researchers have approached it from a purely theoretical perspective (e.g., Calvert, 1985; Crawford & Sobel, 1982), including some that expanded the theoretical analysis to include numerous actors (e.g., Rolfe, 2012; Siegel, 2009). Others built on these theoretical models to run experiments (e.g., Carlson, 2019) occasionally with “real stakes” (e.g., Ahn et al., 2014; Krupnikov et al., 2020; Lupia & McCubbins, 1998). Obviously, researchers have used the workhorse of public opinion studies, the survey, to examine social influence (e.g., Huckfeldt & Sprague, 1995; Mutz, 2006; Sokhey & McClurg, 2012). Increased computing power has allowed for the expanded study of social influence with actors embedded inside larger social networks (e.g., Lazer et al., 2010; Song & Eveland, 2015).

Regardless of approach, the fundamental building block of all of these studies is a dyadic interaction that takes on a similar form. An individual (the sender) provides information to another individual (the receiver). These interactions are at the heart of even the fanciest agent-based model or social-network

*Examining Motivations in Interpersonal Communication* 7

analysis. The analyses that expand beyond the two actors are simply studying the dynamic consequences of *many* of these dyadic interactions happening at once or over time.

Experiments are particularly well suited for this type of study because observational studies are plagued by endogeneity issues that are difficult to resolve (Fowler et al., 2011; Huckfeldt, 2007; Huckfeldt & Mendez, 2008) and because they also lack important information necessary to demonstrate actual social influence. For example, there is some level of individual choice in social influence. Many people actively seek out information via both the news media and their friends – further, they choose to accept or reject the information they receive (Zaller, 1992). On the other hand, individuals do not have *full* control over the flow of information and thus over social influence. In the modern world, the acquisition of information has the feeling of drinking from a fire hose. Messages can essentially sneak in, producing important consequences for citizens’ informational biases (Huckfeldt, 1983; Huckfeldt & Sprague, 1995). While most surveys will miss these nuances (see Carlson & Settle, 2022), even in the most thorough survey, neither these messages nor their source will be reported due to respondents’ lack of awareness they are receiving them.

Carefully designed laboratory experiments can address these obstacles, allowing for analysis at the dyadic level that can form the basis of the study of more complex processes. As we discussed in Section 1, the experimental frameworks we rely on in this Element are incentivized studies based on a Downsian (Downs, 1957) spatial model of political preference and competition. And as we have noted in *this* section, there are many other approaches – including experimental approaches – to addressing questions related to social influence. At the same time, we see several advantages to utilizing a spatial model. First, these designs allow us to easily assign all potential senders, receivers, and candidates a position on a single dimension. Hence, it becomes possible to talk about the “distance” between any pair of citizens or candidates on a continuum – that is, preferences.

Second, we can more easily incorporate concepts of interest that are difficult or impossible to randomly assign in an experiment (let alone to examine in an observational setting). For example, our interest in political expertise in Section 3 necessitates random assignment of political expertise. The Downsian framework allows us to do this: we can examine political expertise by incorporating political uncertainty. Of course, political expertise presupposes uncertainty – if everyone has perfect information about politics, then everyone will be equally expert at politics, making the concept of expertise irrelevant.



Indeed, spatial models of political competition abound in which citizens are more or less informed about the candidates' true positions or the ultimate impact of a proposed policy on the citizen's own welfare (e.g., Baron, 1994; Budge, 1994; Calvert, 1985; Shepsle, 1972). Level of expertise can be represented in various ways, including the proximity of a voter's estimate of a candidate's true position to the actual policy position of the candidate, the variability of a voter's estimate of a candidate's position, or the amount of information a voter has acquired. In this way, we can incorporate not only the heterogeneity in preferences but also variation in expertise.

These advantages are well understood and well studied. One potential issue is that they may also be well understood by experimental *participants*. As discussed in Section 1, the way researchers communicate experiments to participants can lead participants to behave in particular ways (Groenendyk & Krupnikov, 2021). As a result, when a design focuses on preferences and expertise – and, therefore, experimental instructions explicitly mention preferences and expertise – participants might believe that they should primarily consider preferences and expertise in their decision-making within the experiment. Participants concentrating on preferences is especially concerning given that effective interpersonal communication is difficult because of the frequency with which people send biased information (e.g., Carlson, 2019; Ryan, 2011b). Leading participants to prioritize preferences might exacerbate this, which might suggest that interpersonal communication is hopelessly plagued by motivated reasoning and self-interest – *or* it might be that participants act self-interested because they believe that is what the researchers want them to do.

The good news is we can amend traditional group-based, incentivized experimental designs to take into account other potential motivations in political discussion *beyond* the simple partisan motivations of the sender. Doing so is crucial because observational work suggests the effects of interpersonal communication are strongly correlated with participants' motivations (Eveland, 2004). One advantage of previous experiments is that expertise is manipulated by presenting participants with a set of information that is unbiased – this avoids inferential problems caused by partisan differences in information the subject brings to the experiment (Tappin, Pennycook, & Rand, 2020) and allows for clear identification of behavior consistent with partisan motivations. The weakness, however, is that researchers cannot test what consequences result when partisan motivations are mixed with *other* motivations.

Despite the plethora of studies looking at interpersonal communication, one would be hard-pressed to write a literature review of experiments examining



## *Examining Motivations in Interpersonal Communication* 9

how different motivations affect social influence. An exception is Pietryka's (2016) study, which explicitly introduces accuracy motivations to the design – demonstrating that people do not necessarily make more accurate decisions when they are motivated to do so. This finding also demonstrates a broader implication for the study of motivations: if accuracy motivations do not necessarily increase accuracy, the *effects* of motivations may differ from their *intent*. Thus, since these effects cannot be inferred, we can better understand them through experiments that randomly assign different motivations.

The experimental innovations we discuss are similar to Pietryka's as we also randomly assign nonpartisan motivations to participants. Random assignment to motivations is the only way to know that the outcomes we observe are the result of differences in motivations and not some other process we failed to consider (Leeper & Slothuus, 2014). By the end of this Element, we will demonstrate that these types of design choices have implications for our understanding of the effects of interpersonal communication. Introducing other motivations can either amplify or attenuate the bias in both the messages people send and the perceptions of the messages people receive, ultimately shaping how democracy functions. Hence, the efficacy of a discussion depends on which motivations shape that discussion.

### 2.1 Current Approaches and Findings

At the most basic level, researchers conduct interpersonal communication experiments in the hopes of understanding what will happen in versions of the following hypothetical scenario: someone, who we will call Rob, is deciding between two or more options and someone else, who we will call Samara, provides him with some information about the options. Researchers build their treatments around this basic scenario. For example, the information Samara provides Rob may not be accurate either due to Samara's ignorance or deliberate dissembling. Or, Rob may not need Samara's information – Rob may actually already have enough information to make the decision that is best for his interests. Experimenters might *also* vary whether or not Rob specifically requested information from Samara or if Samara was assigned to send information to Rob.

Another key factor in almost all of these experiments is whether Samara and Rob have preferences that are aligned – recall from Section 1 this is important in the Downsian framework. That is, an important experimental treatment is whether Samara is better off if Rob makes the right choice for Rob (i.e., their preferences are aligned) or if Samara is *worse* off if Rob makes the right choice for Rob (i.e., their preferences are at odds). This could be manipulated in several