

How You Ask Matters: Interview Requests as Network Seeds

AshLee Smith
Ph.D. Candidate
University of Minnesota
smit8385@umn.edu

Jane L. Sumner
Assistant Professor
University of Minnesota
jlsumner@umn.edu*

Josef Woldense
Assistant Professor
University of Minnesota
jwoldens@umn.edu*

Abstract: When recruiting interview subjects is the goal, building rapport is conventionally heralded as the superior method. Cold-emails, in contrast, are often dismissed as inferior for their low response rate. Our study suggests that this stance is mistaken. When it is elites who are to serve as interview subjects, we argue that cold-emails can yield tremendous benefits that have thus far been overlooked. More specifically, we posit that when paired with network effects, which are rooted in the linkages among elites, cold-emails can outperform the standard but costly interview solicitation method of building rapport with subjects. In a series of experiments and simulations, we show that small changes to the wording of cold-emails translates into greater network coverage, thereby offering researchers a richer set of insights from their interview subjects.

* Corresponding authors.

Introduction

Interviews stand alone in their capacity to inform researchers about situations from first-hand accounts of people on the ground. From learning about Medicaid from those reliant upon the system (Michener, 2018) to learning about nuclear proliferation from elites responsible for crafting policy (Bell, 2021), interviews have proven invaluable in fortifying knowledge for political scientists. They can be useful for theory building, hypothesis testing, or even informing priors in statistical models (Nunes, n.d.).

Yet interviewing can come with enormous barriers to entry. Unlike many other methods, the barriers are not primarily in the form of training or equipment, but rather in social capital. Conducting interviews requires the researcher to first successfully recruit interview participants, which is no small feat. This is especially true when setting out to interview elites, who may be difficult to reach and may value their time too highly to agree to participate.

Although there is ample material on the nuances of conducting interviews (Mosley, 2013), the critical prior step of recruiting elites receives comparatively little attention. Consider, for instance, Lynch (2013) who while acknowledging that researchers “may not feel important enough to commandeer the time and attention of high-status individuals like government ministers, elected officials, or CEOs,” nonetheless concludes that in her experience and the experience of others she cites, “such concerns about gaining access [are] unfounded” (p. 43). This, of course, overlooks the social capital and experience that scholars like Lynch can leverage. Absent such a professional pedigree, access is a considerable problem, one that can be so daunting as to dissuade the researcher entirely.

The purpose of this study is to examine the issue of access, but with the explicit assumption that the researcher lacks social capital. More specifically, our focus is on “cold-emails” -- the act of sending unsolicited emails to elites in the hope of obtaining an interview -- which is becoming increasingly the most readily available recruitment strategy to researchers. The problem is that it is also seen as the least desirable method because it suffers from low response rates. Indeed, according to conventional

wisdom, the best method is building rapport with potential interview subjects, followed by referrals in the form of snowball sampling and trailed by cold-emails as the worst of the three approaches. There is a good reason for this hierarchy. If the objective is to maximize the success rate for a given interview request, rapport certainly is the best approach. By this same metric, it is also true that cold-emails are likely to perform far worse than the other two strategies. Perhaps not surprising, then, that cold-emails are viewed as inferior and frequently dismissed.

In the coming pages, we shall argue that such sentiments are misplaced. Not only can cold-emails offer tremendous benefits to researchers, but under certain conditions, they can even outperform rapport as a recruitment tool. In making these claims, we propose an alternative conceptualization that redefines the purpose of the aforementioned recruiting methods. Rather than maximizing the response rate for a given interview request, we stress that in the context of elites, the goal is maximizing access into a network. Where the relevant interview population is interlinked to form a network, as is the case with elites, the researcher's task is to access that network and then subsequently explore it. The latter can be accomplished via referrals where snowball sampling enables the researcher to penetrate the network ever more deeply. Access, in contrast, is done through rapport and cold-emails. This first step of seeking access is tantamount to sowing seeds across the network and conditional on those seeds sprouting, access gives way to referral cascades.

It is this network amplification effect that we argue is missing from current debates around the value of cold-emails. As we will show in the paper, once this effect is included, there are two distinct benefits that emerge. First, the amplification effect translates relatively minor increases in the response rate of cold-emails into significant gains in network coverage. A difference of only six percentage points in the initial response rate, for instance, can grow almost seven fold in terms of network coverage. This finding is important because it suggests that rapport is likely the wrong benchmark when evaluating the response rate of cold-emails. Second, we show that when the quantity of emails offset the lower response rate, cold-emails outperform rapport in the

diversity of subjects that are recruited. The reason for this is because rapport is limited in both the quantity and geographic dispersion of the seeds it can sow. Consequently, it traps the researcher in certain neighborhoods of the network. That is to say, without sufficient referral cascades, rapport produces a more parochial view than that produced by cold-emails.

Taken together, once we properly frame email solicitation in network terms, we see that far from an inferior method, it has great potential to serve researchers in their efforts to secure interviews. Rather than focusing on low response rate, cold-emails should be praised for their capacity to seed an interview network and, therefore, allow the researcher to permeate and explore more of the elite network than they might otherwise be able to. As such, we argue that attention to cold-emails is crucial, since even a small increase in response rate can lead to a significant increase in elite network coverage.

In the next section, we provide a brief overview of the existing research and guidance about interviews, and compare interviews with their closest cousin: surveys. Then we describe and discuss three empirical analyses: a conjoint experiment, a network simulation, and a firm experiment. Each of these analyses demonstrates a different aspect of our claim. We use the conjoint experiment (and an in-process firm experiment) to demonstrate that 'wordsmithing' does, in fact, increase response rates, even if that increase is modest. We use those results to inform the parameters of a network simulation, which we use to demonstrate the network effects that translate increases in response rates into cascades of referrals, as well as demonstrate how seeding interviews throughout a network produces a more representative network than one seeded by rapport or referrals. Third, we describe the firm experiment, which is still in progress, that offers empirical validity for the assumption that elites form a network and to provide some empirical evidence for the network simulations.

A Brief Overview of the Method of Interviewing

Interviews are a common method of inquiry in the social sciences as they can be used to “yield rich insights into people’s biographies, experiences, opinions, values, aspirations, attitudes and feelings” (Cassell 2005; May, 2001, p. 120). Interviews are regarded as the gold standard method used in qualitative research as they are an important, and often a fundamental tool for making sense of a phenomena of interest. They can be used to collect information from individuals about their own behaviors, beliefs, or opinions (Rubin and Rubin 2011). Through interviews, one can get to the core of a subject of interest’s processes and attitudes, gather background information, or delve into an expert’s knowledge. Rubin and Rubin (2011) note that interviews “are especially good at describing social and political processes, that is, how and why things change” (Page #). Soss (2006) writes that interviews, “permitted [him] to treat client statements as more than a series of discrete verbal reports to be coded, each in its own right, and then correlated with one another. It allowed [him] to pursue the meanings of specific statements by locating them within a broader web of narratives, explanations, telling omissions, and nonverbal cues” (128). If we are going to use this method of inquiry to pursue and understand puzzling phenomena, the researcher must first make personal contact with those they wish to understand.

Thus, to conduct interviews the researcher needs access. Access to people, especially people in power, varies considerably and can be the hardest challenges to overcome for researchers (Jackall 1988; Ostrander 1993). Many other scholars have written tips to overcome this challenge, but they all focus on the same thing, *rapport* (Burnham et al., 2004; Dexter, 1970; Lilleker, 2003; Harvy, 2011; Aberbach and Rockman, 2002; McGrath et al.2019). Rapport is “about trust – enabling the participant to feel comfortable in opening up to you” (King & Horrocks 2010, p. 48). Ostrander (1993) argues that rapport is built up over time and researchers should attempt to build a rapport with their subjects from the moment they first contact them. Unfortunately, it takes a long time to build rapport and such an approach is a function of the resources and the various identities that the researcher may have. Glasius et al. (2018)

acknowledge this clearly, noting that “our choices [about how to obtain interviews] are very much constrained by age, gender, and position, but each of us can work with what they have”. Yet even working with ‘what you have’ means that a strong focus on rapport can disadvantage scholars approaching the situation with different identities that may be seen as less credible given the situation, or who may have fewer resources to dedicate to the pursuit of rapport.

In lieu of having a personal connection with potential research participants prior to an interview request, many researchers rely on snowballing - a recruitment strategy that requires researchers rely on their participants' social networks (Creswell 1998). However, this strategy presupposes that the researcher already has some connections and rapport with their subjects of interest. Cold-emails are often derided as an inferior method, for those unwilling to put in the time to ‘pound the pavement, hit the bars, and use their networks. Consequently, there is minimal discussion of it as a method, or guidance on how to do it well.

Interviews v. Surveys

In contrast with interview solicitation, there has been a great deal of research done on sampling for surveys, as well as on the importance of question wording in surveys. There is research on whether online platforms yield more or less representative samples (e.g., Coppock and McClellan 2019, Huff and Tingley 2015, Mullinix et al 2015, Krupnikov and Levine 2014), on how the wording of questions matters in shaping the kinds of responses a researcher gets (e.g., Hiscox 2006, Schaeffer and Presser 2003), and on the consequences of survey non-response (e.g., Bailey 2020).

Yet although interviews and surveys share some things in common -- notably that they ask a group of people to provide information to the researcher -- their objectives are different. To put a fine point on it, interviews are not just small-N surveys, and surveys are not large-N interviews. Surveys typically aim to estimate parameters of interest along with uncertainty bounds -- for instance, the proportion of the population

that supports some policy, or the effect of one variable on another. Interviews, on the other hand, aim to inform the researcher about an entire world.

One crucial difference between elite interviews and surveys is that elites tend to be connected to one another, which gives rise to a particular network structure. In surveys, the networks between participants is not important,¹ whereas it is a defining feature in elite interviews. It is because of this network structure that referrals and snowball sampling are key to broadening the researcher's interview base and therefore expanding the part of the world they are able to see and study.

In soliciting elite interviews, the researcher, therefore, has two challenges: first, accessing that network, and second, using snowballing and referrals to cover it. It is here that building rapport becomes a less attractive strategy. While building rapport may yield a higher response rate for a given solicitation request, it is extremely limited in the number of network access points it affords the researcher. After all, building rapport can only be done successfully when the number of people is small.

Cold-emails, on the other hand, are able to cast a much wider net and thus target multiple entry points into the network simultaneously. Even in the face of low response rates, such broad targeting provides the researcher with a more diverse pool from which to perform further snowball sampling. That is to say, cold-emails allow researchers to set more seeds across the network, which translates into more opportunities for snowball sampling to recruit more subjects, thereby increasing network coverage.

What this means is that, although the probability of response from each email is low, each seed brings with it enormous promise for broad network coverage. That suggests that cold-emails need to be approached and composed with care and deliberation, a practice we call 'wordsmithing'. If that care translates into even a small increase in the probability of response, we argue that it can, potentially, lead to a substantial increase in the network coverage.

Taken together, this suggests a few testable propositions:

¹ Under certain circumstances, and especially if the survey has an embedded experiment, networks among participants may actually be harmful because they may violate the assumption of independence between observations..

(1) cold-emails that are written more intentionally, especially with an eye to the incentives of the potential recipient to respond and agree, should have a higher response rate than those that do not.

(2) cold-emails that are written more intentionally should result in broader network coverage than those that are written less intentionally.

In the next three sections, we assess these propositions using two methods (online experiment and network simulations), and then explain and provide some preliminary results for a third (experiment with corporate elites).

Conjoint Experiment

First, we empirically validate the key assumption underlying this project: that the wording of emails matters, and that emails that pay more attention to their wording and presentation have a higher success rate than less careful emails. To figure out if this is supported by evidence, we conducted a conjoint experiment on the survey platform Lucid² (Hainmueller, Hopkins, and Yamamoto, 2014).

For this experiment, we primed participants to think of the industry they are most familiar with³, then asked them to put themselves in the position of someone in a managerial position in that industry. We gave participants a pair of emails that were identical except for four aspects: our appeal to their position as an authority (*authority*), the amount of time proposed for the interview (*time*), a description of the confidentiality of the interview (*confidentiality*), and the description of the problem they'd be helping to solve, an appeal to their authority (*problem*). Participants were then asked to rate each email on a scale of 0 to 100, representing how likely it is they would respond, and then to pick which email they are most likely to respond to.

Generally, each email varied in its 'wordsmithing': the amount of care put into crafting the particular appeal, ranging from a low level (little thought of the recipient's incentives or position) to a high level (considerable effort and thought about the

² For more on Lucid, see Coppock and McClellan 2019.

³ Specifically we asked them what industry they are most familiar with, and then we asked them to briefly describe what tasks a manager in that industry was typically responsible for.

recipient's incentives). We chose each of these aspects by considering the perspective of the hypothetical recipient, and tailoring the email to speak to their concerns and incentives.⁴ For *authority*, we varied whether we recognized the recipient as an important figure with a valuable perspective, or not, under the assumption that people are more likely to want to help when they understand they are important. We varied *time* because we assumed our recipients would value their time highly, and would like to know how much time to set aside, and not give much away. *Confidentiality* varied between a more and less detailed statement of interview confidentiality, laboring under the assumption that people speaking on behalf of a company would be particularly concerned about appearances, liability, and the potential for information to be compromised. Finally, we varied the *problem* we were asking them to help with, from the narrow (helping a researcher understand a question) to the universal (correcting public misperceptions). Each of these four aspects had between two and four treatments that were independently chosen at random⁵ (see Appendix A for full text of skeleton email and treatments).

For our purposes, we're interested in assessing the effects of wordsmithing generally, rather than the effect of any particular treatment. Our experimental variations reflect our hypothetical audience, which is not generalizable. That is, most researchers do not need to know if telling a manager of a company an interview will take 15 minutes is better than not giving them a time estimate at all. What is generalizable is whether our conscious efforts to take the perspective of our recipient and carefully consider their incentives and position resulted in a higher response rate. As such, our strategy is to pair two emails that differ in the types and amount of wordsmithing they carry and then observe whether there is a measurable difference in the kinds of emails that are

⁴ These assumptions drew on experience soliciting interviews used in Sumner (n.d.).

⁵ We use Javascript (within Qualtrics) to randomly choose and order pairs of treatments, informed by the work of Strezhnev et al (2013).

chosen. At its simplest, we conjecture that, all else equal, a more wordsmithed email is more likely to be selected.⁶

To this end, Figure 1 offers descriptive evidence that seems to support our claim. Recall that we performed wordsmithing along four dimensions: authority, time, confidentiality and problem. Collectively, the four dimensions constitute our treatment regimen so that a given email can be described as $e_{1,0,0,1}$ where 1 means a treatment -- any of them -- is present and 0 means treatment is absent and the aspect was blank.⁷ The subscript [1,0,0,1] is thus read as: authority = 1, time = 0, confidentiality = 0, problem = 1. Building on this, we can determine how different a pair of emails are by subtracting their treatment vectors. For example, suppose email A has treatment [1,0,0,1] and email B has [1,1,1,1] such that subtraction yields $[1,0,0,1] - [1,1,1,1] = [0,-1,-1,0]$ with the resulting sum being -2. Here, email B carries two more treatments than A, which gives us the distance score for this specific pair. The largest distance is ± 4 , where negative means that B has more treatments and positive means that A has more treatments.

With this, we can ask a simple question: Given, for instance, a distance score of -4, what is the probability that email B is chosen? Or stated more generally, given a distance score of $\pm x$, what is the probability that the 'correct' email is chosen? Figure 1 displays the results of this exercise. Although left out for convenience, the values on the x-axis should be thought of as having a \pm sign in front of them. Where the distance score is 0, Figure 1 shows the baseline probability for selecting A or B, which, interestingly, was exactly the same (selected 33 times out of 95 appearances). The key result, however, is the upward trend in the direction we would expect it to go. When moving from a distance score of 0 to ± 4 , the corresponding increase for a given email to

⁶ For this exercise, the scores given to each email appear less useful than the forced choice because personal characteristics -- which we do not control for at this time -- play a large role in determining someone's baseline scores.

⁷ For the time aspect, 0 represents null as well as 45 minutes, while treatment is 15 minutes. This is because the comments from participants clearly showed that the longer interview time was looked at very negatively, which is the opposite of what we aim for with wordsmithing. This choice of coding does not affect our substantive findings.

be selected is 13% (from 35% to 48%). A greater gap in wordsmithing results in a higher probability of selection.

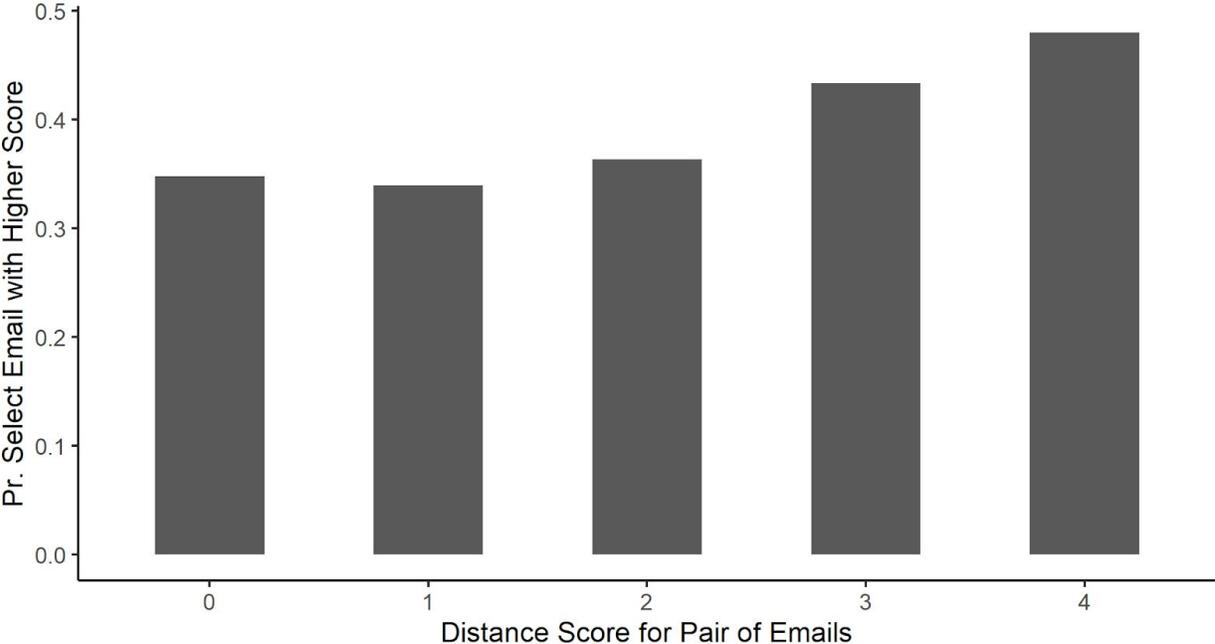


Figure 1: The probability an email is selected based on the distance score between email pairs.

Figure 2 further breaks down the combination of treatments and their success rate. Each dot represents the probability that an email was selected, given the treatment vector represented by black squares and blank spaces on the left side. For instance, the top row shows that treatment [0,1,1,1] yielded a 50% success rate, since it appeared 53

times in the survey and was selected 27 times. The emails that perform the worst, near the bottom of the plot, are the ones with the fewest treatments. The relationship is not perfect, because different treatments have different strengths, and the treatments themselves are specific to this application. There are some vectors with two that outperform those with three, for instance. Yet as Figure 2 illustrates, as the number of treatments increases generally, so too generally does the probability of success.

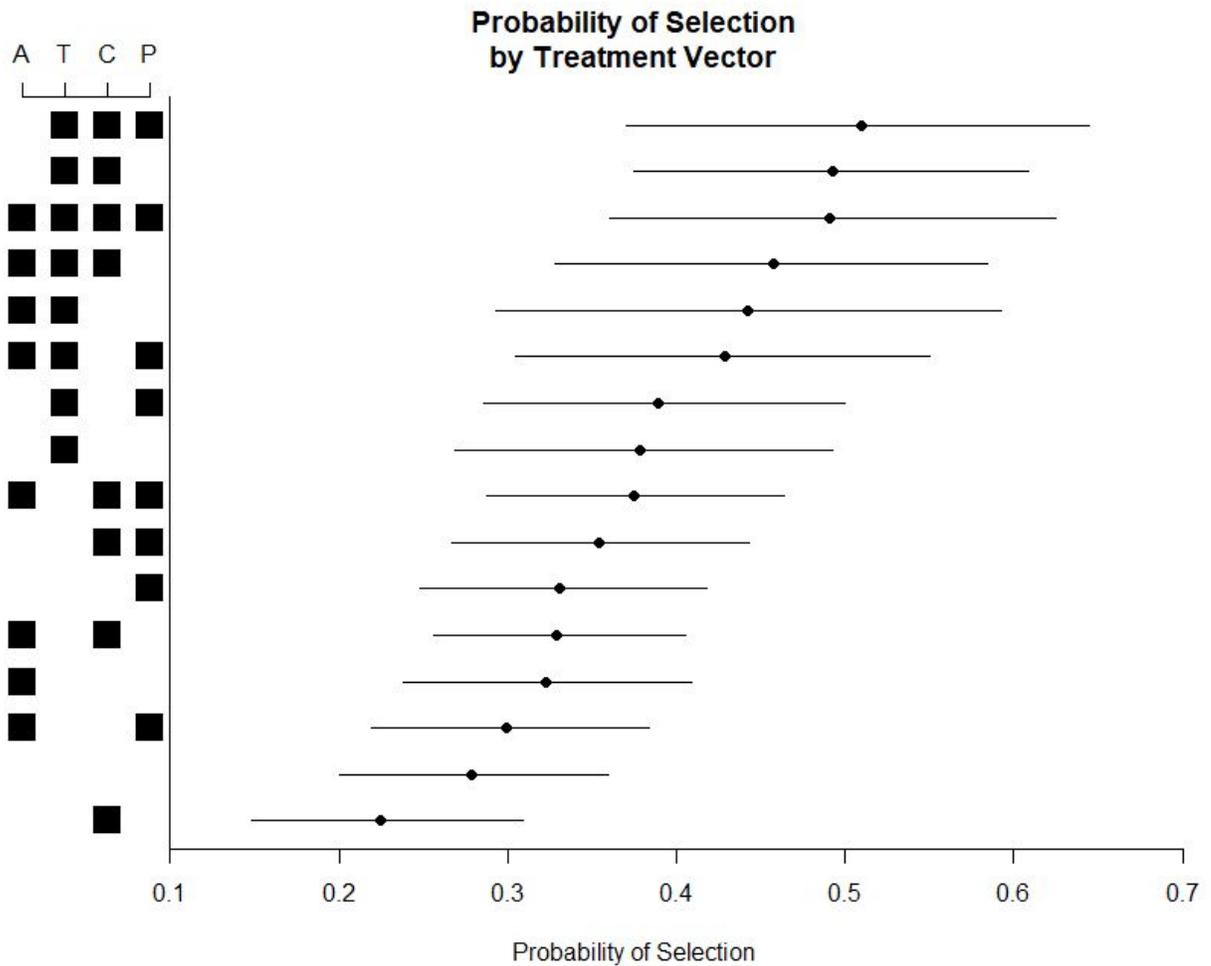


Figure 2: The probability that any email is chosen, broken down by treatment vector. The configuration of squares on the left side indicates the treatment vector, where black squares are a 1 and blanks are a 0. Confidence intervals were simulated using 10,000 interactions of non-parametric bootstrap.

Figure 3 tells a similar story by looking at the sum of the treatment vector -- the total number of wordsmithed treatments, without regard to which treatments are assigned -- and the corresponding probability of being selected. In general the more wordsmithed the email, meaning the more treatments and therefore the more thought, the higher the probability the email was selected when paired against another set of treatments. Because this figure combines treatment vectors, it reflects the same insight as Figure 2, that not all treatments are equally effective and therefore probability is not necessarily increasing in number of treatments. Yet the general point stands: the more thoughtfulness goes into the email, the more likely the email is to be selected.

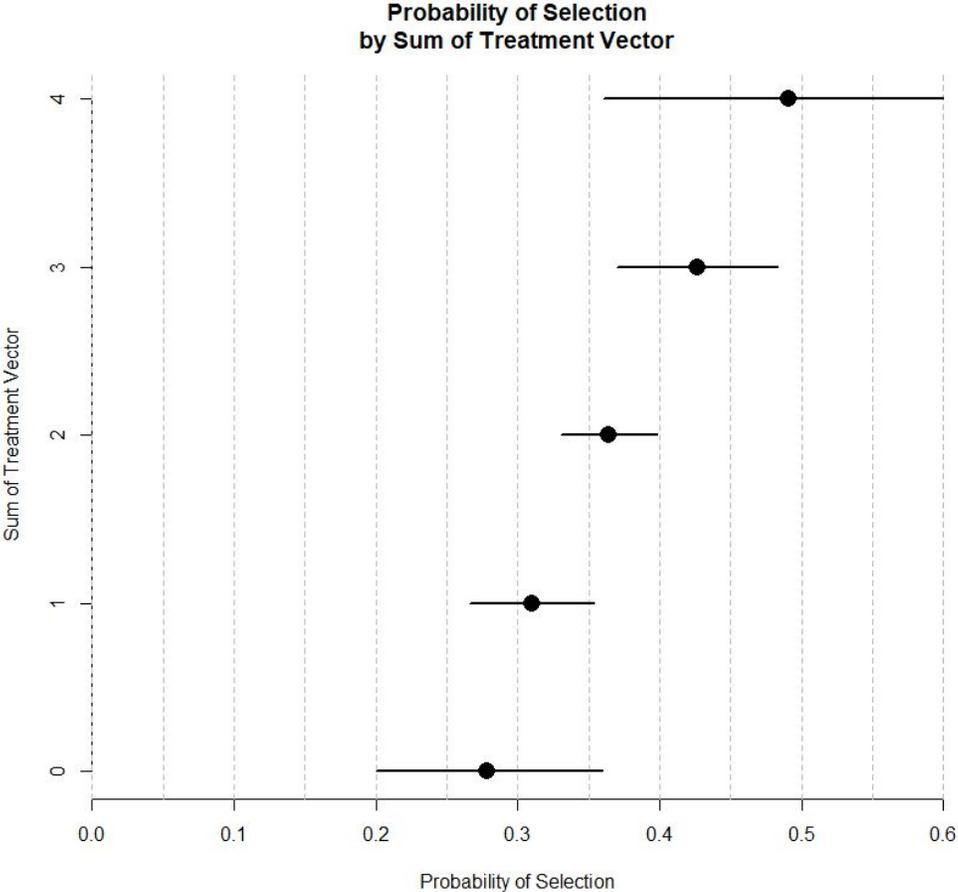


Figure 3: Probability of selection by total number of treatments. Confidence intervals were simulated using 10,000 iterations of non-parametric bootstrap.

This all serves as a proof of concept. Among an internet sample faced with a hypothetical situation, cold-emails that are more thoughtful and “wordsmithed” in subtle ways, even if entirely the same otherwise, have a slightly higher success rate. In the next section we explain why even this small increase in response rate is important and can lead cold-emails to outperform rapport on network coverage.

Network Simulations

Having established the key assumption of our project -- that small changes and attention to writing cold-emails can increase response rate -- there is an obvious follow-up question: does it matter? After all, even if wordsmithing increases the response rate in paired comparisons, two things are also true. First, the probability of response is still significantly lower than we might expect for rapport even in the best of our experimental outcomes. And second, most elites are making a different choice than our experimental participants did -- they are not viewing two emails and deciding which one they prefer, they are viewing one email and deciding whether to respond or not. This suggests our experimental probabilities may be an overestimate, and they are not high to begin with. If that is the case, then, what does it matter that wordsmithing increases a very small probability by a small amount?

To answer these questions, we use network simulations to reveal the logic by which first order responses (i.e., solicitation through email or rapport) translate into second order responses (referrals). That is to say, we examine how the network in which elites are embedded maps initial interview requests onto subsequent referral cascades. Figure 4 offers an illustration of this process where the nodes in red are the seeds in period one while the nodes in yellow represent the succeeding interview subjects recruited through referrals. As displayed in Figure 4, referral cascades travel along the linkages among the nodes so that elites with many connections can serve as ‘super spreaders’. Since the question at hand concerns the seeding strategy, the comparison of interest is between wordsmithed emails vs. plain emails and wordsmithed emails vs. rapport. In the former, we are interested in ascertaining the

downstream effects of a marginal increase in the response rate. For the second comparison, we want to identify if and under what condition cold-emails outperform rapport.

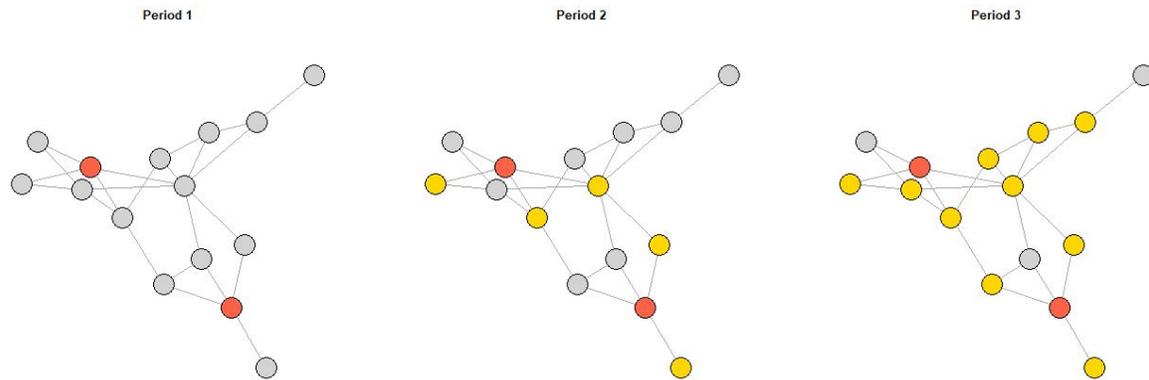


Figure 4: Network seeds and referral cascades. Nodes in red are seeds and yellow represent the interview subjects recruited through referrals.

We assess these analyses along three outcome measures that we call -- *node reach*, *community reach* and *community match*. Node reach is the proportion of nodes in the network that have been successfully recruited as interview subjects, and thus speaks to the number of elites one is able to reach, irrespective of the network structure. The other two measures are more technical, but address an intuitive idea. In recruiting elites, we are not only interested in the number of interview subjects one secures, but also in how diverse they are in their perspective. As research on social networks has amply demonstrated, when people are densely interconnected, they tend to hold similar perspectives. This is in part because the same information easily passes around tightly knit sub-communities.⁸ All else equal, a sample with subjects from different

⁸ Perhaps the most prominent example in the contemporary moment is the political polarization in the US with its distinct information ecosystems, sub-communities that are densely interconnected and regularly share the same information amongst each other.

sub-communities is more valuable to a researcher than one that draws its participants from just one sub-community.

Our analysis captures this idea by running a community detection algorithm that identifies sub-communities, or sub-groups, within the network based on the density of links among nodes. Those sub-groups are then used as the basis for the *community reach* and *match* measure. *Community reach* simply counts the number of sub-groups that are represented in the sample. For this measure, securing just one person from a given sub-group is sufficient for it to be counted as represented. *Community match*, in contrast, is more demanding and requires the sample to match the network's sub-group distribution. Here, a sample is considered least representative when all of its subjects are drawn from the network's smallest sub-group. Placed on a spectrum, we can think of one end of the continuum representing a sample that perfectly mirrors the network's sub-group distribution while on the other end of the spectrum lies a sample which is dominated by the network's smallest sub-group.

Given these outcome measures, we begin by producing a small-world network⁹ with 200 hundred nodes. In period 1, the seeds are set across the network with a response rate of 0.01, 0.08 and 0.8 percent for plain, wordsmithed and rapport respectively. The initial seed size (i.e., number of nodes) is set to 80 for the emails and 8 for rapport. The expected value for wordsmithed emails and rapport is therefore the same, which enables us to hone in on the comparison of interest. Lastly, whereas the seeds are distributed randomly for emails, in the case of rapport, the selected nodes are connected to each other -- thus reflecting the spatial proximity of the subjects one builds rapport with. Once period 1 is completed, period 2 commences the referral cascade where a neighboring node accepts an interview request with a probability of 0.3. All subsequent periods repeat this process with the same probability of acceptance.

Figure 5 shows the results for each of the three outcome measures. Focusing first on the comparison between 'plain' and wordsmithed emails, we can see a striking difference between the two in all three outcomes, despite only a small difference in the

⁹ Produced by Duncan and Strogatz, this is the network structure that gives rise to the well known finding that we are all, on average, only six degrees away from a famous person.

initial response rate. A marginal change in the seed response rate, from 0.01 to 0.08 percent, compounds rapidly with the passing of each period. Consider, for instance, the difference in *node reach*. The difference between the two cold-email groups is only two percent in period one, but then grows to 16 percent in period three and 34 percent in period five. The difference is even more pronounced for the other two outcome measures. The difference between the two groups starts at around 30 percent in period 1 and grows to over 50 percent by period five. These results clearly show that increasing a very small probability by a small amount, as wordsmithing does, carries a tremendous payoff.

Figure 5 also shows when wordsmithing outperforms rapport. As mentioned above, the parameters were set such that the expected outcome for the initial seeds are the same in the first period -- meaning that although the response rate is higher for rapport and much lower for the cold-email, the cold-email sows significantly more seeds. We can see this in the *node reach* panel where the two approaches have nearly identical scores in the first period. Yet the two differ substantially in the other outcomes, and particularly in the *community match* measure. Although the gap between them decreases over time, in period one, cold-email outperforms rapport by 41 percent and remains at 8 percent in period three. In other words, without sufficient referral cascades, rapport struggles to match wordsmithed emails in the diversity of its sample. This means that the researcher who deploys wordsmithed emails is likely to encounter much more diverse voices early in their interview gathering process, and to reveal more of the network through subsequent referrals. The consequences of this can be immense not only for coverage but also for substance and depth, since, in practice, the questions one asks of interview subjects is itself informed by the insights of previous interviews.

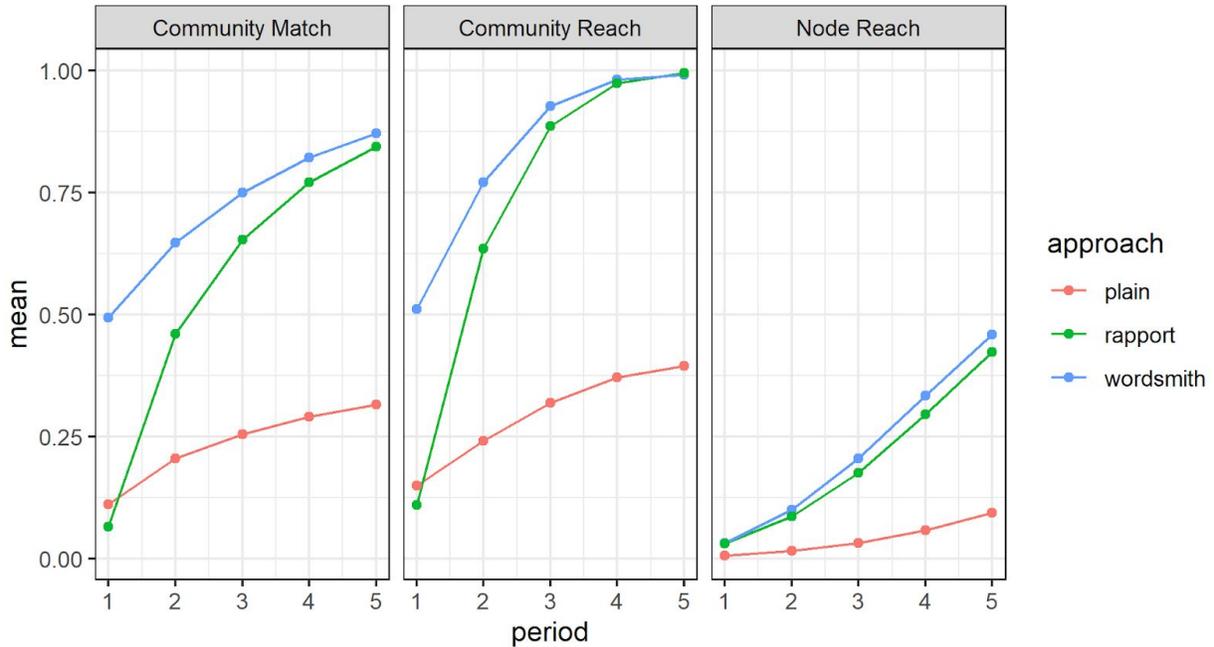


Figure 5: Results for the three outcome measures of the network simulation.

These network simulations support our initial contentions: first, that cold-emails often outperform rapport for network coverage, and second, that even a small increase in a small probability can lead to a significant increase in network coverage. This suggests not only that cold-emails have been unfairly maligned, but also that it is worth the time and effort to carefully craft those emails. If the researcher can successfully sow many seeds across a network, even if only a relatively small fraction of them ‘sprout’, the subsequent referral cascades can result in at least as many nodes as one can access through rapport, but also can result in the researcher accessing a far more diverse sample than they otherwise would.

In the next section, we explain the steps we have taken so far to put our theoretical claims to test in the real world, using an experiment with both low-wordsmithed and high-wordsmithed cold-emails to sow seeds in a large, elite network: publicly-traded US companies.

Firm Experiment (in-progress)

Finally, to both demonstrate that greater wordsmithing of emails increases response rate, and to demonstrate and empirically validate our claims about elite networks, we are in the process of conducting an experiment with publicly-traded companies in the United States. For many questions in political economy, companies are key strategic actors -- for instance, even though foreign direct investment (FDI) is often studied at the aggregate level (as flows or stocks), they often reflect theories implicitly about corporate decision-making (see, e.g., Kerner 2014). Interviews with corporate elites are especially useful for understanding the causal mechanisms in these theories, as well as vital background information on the phenomena of interest. Yet corporate elites, like any elite, can be difficult to access, especially without the time or clout to build rapport.

Given their centrality to so many political economy questions, we choose to test this theory by trying to recruit corporate elites to be interviewed about their companies' political concerns and advocacy strategies.¹⁰ We randomly sampled 400 companies (11.1%) from a sample of 3,572 companies that are or have been publicly-traded on U.S. stock exchanges since 1994.¹¹ We were able to find publicly-available contact emails¹² for 377 of the companies (94.0%).¹³ Each company was randomly assigned to receive either a wordsmithed treatment or a non-wordsmithed control (see Appendix B). Both began with the same skeleton email, but the treatment added information informed by the findings of our conjoint experiment and our knowledge of the incentives of the recipients, while the control centered the needs and perspectives of the researcher. We successfully sent 365 emails.

¹⁰ This is the topic of Sumner (n.d.). We are recruiting interviews for that project.

¹¹ We sample from a list of publicly-traded companies because it is straightforward to quantify and visualize the overall population, which is not true of companies as a whole.

¹² The majority of these public contact email addresses are to investor relations or media and public relations. To find the emails, we first looked at the company's website, then at press releases, and finally at Securities and Exchange Commission (SEC) documents. In the few cases where none of these yielded an email address, we stopped looking, figuring that companies that did not want to be contacted were especially unlikely to respond.

¹³ Of the 24 companies for which we could not find contact emails, five (20.8%) had either been delisted, were no longer operating, or were "blank check", trust, or shell companies..

Those emails were sent July 7, 2020. For each company, we are tracking whether the company responds in any fashion and whether they decline or agree to be interviewed. We conclude every successfully-conducted interview by asking for referrals. We will continue this process until the end of August. These data will allow us to test three things. First, we can test whether the wordsmithed emails lead to a higher response rate than the non-wordsmithed emails, controlling for a company's industry, location, and size. Second, we can test whether they have a higher rate of agreement conditional upon response. Third, we can estimate the true referral probability and assess whether the wordsmithed emails ultimately lead to broader network coverage than the non-wordsmithed emails.¹⁴

Conclusion

Interviews are an especially valuable source of information for political scientists, yet they can be among one of the most inaccessible. The conventional wisdom about recruiting elite interviews is that researchers need to build rapport or have connections. This is a solution that is implausible for many and that can replicate and exacerbate many existing inequalities within the discipline. In this paper, we mount a defense of the cold-email -- emailing someone unknown, out of the blue, and asking for an interview. Often denounced as a decidedly subpar recruitment strategy, we argue that cold-emails, despite having a lower probability of response than rapport or connections, can allow the researcher to seed the elite network more broadly and therefore, with the addition of referrals and snowball sampling, actually cover more of the network than they would by using rapport.

In this paper, we demonstrate evidence for our key assumption -- that subtle, thoughtful additions to an email can increase the probability of response. We do this using a conjoint experiment on the platform Lucid. Then we present the key contribution of our paper. We demonstrate two things. First, that cold-emails fare better than rapport

¹⁴ The researchers have no rapport in this situation, so we cannot plausibly test this against the rapport condition, and as such this empirical evaluation focuses strictly on the comparison of wordsmithed v non-wordsmithed emails.

at covering elite networks. And second, that small increases in the probability of response for a cold-email can generate especially large returns in network coverage. Finally, we lay out the research design for an empirical test of that theory, in which we conduct an experiment on publicly-traded companies to determine whether a more-wordsmithed email actually has a higher response probability, and whether the network coverage we theorize actually obtains.

In this study, we do *not* assess or identify the factors that enabled us to increase response rates in our specific study (i.e. unpack wordsmithing). It is important to recognize that those factors are likely not generalizable. Different contexts will require different wordsmithing, since different audiences have different incentives and concerns. Hence, the purpose of our paper is not a set of prescriptions for others to follow. Instead, we merely wish to impress on our readers that it is worth their time to carefully think about their audience and, having done so, wordsmith accordingly. When possible, researchers could even conduct small experiments where they send different email scripts to a subsample of respondents and settle on the one that shows the most promising results.

In conclusion, we hope our research can accomplish a few things. First, we hope our finding that rapport is not necessary -- and may even be inferior to cold-emails -- encourages more researchers to pursue interviewing as a method, rather than be dissuaded by their lack of connections or social capital. Interviews are among the most valuable ways to learn first-hand about a phenomena for background or theory-building, and can stand on their own or supplement statistical analysis for hypothesis testing. Second, although our findings cannot instruct someone on *how* to craft an email, we hope to convince researchers to put thought and effort into crafting their email solicitations. Even a small increase in a small probability of response can compound and lead to large differences in network coverage, and greater network coverage ultimately leads to more nuanced understanding.

Citations

Bailey, Michael A. "Designing Surveys to Account for Non-Ignorable Non-Response." *Working paper* (2020).

Bell, Mark S. *Nuclear Reactions: How Nuclear-Armed States Behave*. Cornell University Press, 2021.

Cassell, Catherine. "Creating the interviewer: identity work in the management research process." *Qualitative research* 5, no. 2 (2005): 167-179.

Coppock, Alexander, and Oliver A. McClellan. "Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents." *Research & Politics* 6, no. 1 (2019): 2053168018822174.

Creswell, John W. "Qualitative research and research design: Choosing among five traditions." *London: Thousand Oaks* (1998).

Gladius, Marlies, Meta de Lange, Jos Bartman, Emanuela Dalmaso, Aofei Lv, Adele Del Sordi, Marcus Michaelsen, and Kris Ruijgrok. "Building and Maintaining Relations in the Field." In *Research, Ethics and Risk in the Authoritarian Field*, pp. 53-76. Palgrave Macmillan, Cham, 2018.

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. "Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments." *Political analysis* 22, no. 1 (2014): 1-30.

Hiscox, Michael J. "Through a glass and darkly: Attitudes toward international trade and the curious effects of issue framing." *International Organization* (2006): 755-780.

Huff, Connor, and Dustin Tingley. "'Who are these people?'" Evaluating the demographic characteristics and political preferences of MTurk survey respondents." *Research & Politics* 2, no. 3 (2015): 2053168015604648.

Lynch, Julia F. "Aligning sampling strategies with analytic goals." *Interview research in political science* (2013): 31-44.

Kerner, Andrew. "What we talk about when we talk about foreign direct investment." *International Studies Quarterly* 58, no. 4 (2014): 804-815.

Krupnikov, Yanna, and Adam Seth Levine. "Cross-sample comparisons and external validity." *Journal of Experimental Political Science* 1, no. 1 (2014): 59.

May, T. (2001) *Social research : issues, methods and process*. Buckingham : Open University Press.

McGrath, Cormac, Per J. Palmgren, and Matilda Liljedahl. "Twelve tips for conducting qualitative research interviews." *Medical teacher* 41, no. 9 (2019): 1002-1006.

Michener, Jamila. *Fragmented democracy: Medicaid, federalism, and unequal politics*. Cambridge University Press, 2018.

Mosley, Layna, ed. *Interview research in political science*. Cornell University Press, 2013.

Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman, and Jeremy Freese. "The generalizability of survey experiments." *Journal of Experimental Political Science* 2, no. 2 (2015): 109-138.

Nunes, Felipe. "A Two-Part Bayesian Model with Elicited Priors to Analyze Longitudinal Government Expenditures in Latin America." PhD diss., UCLA, 2013.

Rubin, Herbert J., and Irene S. Rubin. *Qualitative interviewing: The art of hearing data*. sage, 2011.

Schaeffer, Nora Cate, and Stanley Presser. "The science of asking questions." *Annual review of sociology* 29 (2003).

Strezhnev, Anton, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. "Conjoint Survey Design Tool: Software Manual." 2013: 69.

Soss, Joe. "Talking our way to meaningful explanations: A practice-centered view of interviewing for interpretive research." *Interpretation and method: Empirical research methods and the interpretive turn* (2006): 127-149.

Sumner, Jane Lawrence. *The Cost of Doing Politics: Corporations, Consumers, and the Consequences of Influence Seeking*. (Revise and resubmit at Cambridge University Press, Business and Public Policy Series).

APPENDIX A: Emails and Treatments for Online Experiment

Skeleton email:

I am an Assistant Professor of Political Science at the University of Minnesota. I am planning to conduct interviews with high-level managers like yourself **[PROBLEM]**. As you know, companies face a lot of political challenges that are poorly understood by the public. I am trying to capture your perspective on the political landscape that your company faces. **[AUTHORITY]**.

[TIME]. Your responses to the questions will be kept confidential. **[CONFIDENTIALITY]**

If you are willing to participate, please suggest a day and time that suits you and I'll do my best to be available. The interview will be conducted by phone or video chat. If you have any questions please do not hesitate to ask.

Please feel free to reach out at your convenience.

I look forward to hearing from you.

Problem:

1. Null
2. "to help me understand the political challenges faced by companies in your industry."
3. "to help people understand the political challenges faced by companies in your industry."
4. "to correct misunderstandings about the political challenges faced by companies in your industry".

Authority:

1. Null
2. "As a manager in your industry, you are in an ideal position to give us valuable first-hand information from your own perspective."

Time:

1. Null
2. "I expect that the interview will take no longer than 15 minutes of your time."
3. "I expect that the interview will take no longer than 45 minutes of your time."

Confidentiality:

1. Null
2. "Each interview will be assigned a code to help ensure that personal identifiers are not revealed during the analysis and write up of findings."

APPENDIX B: Treatment and Control Email for Firm Experiment

TREATMENT:

My name is Jane Sumner and I am an Assistant Professor of Political Science at the University of Minnesota, Twin Cities. As I am sure you know, companies face a lot of political challenges that are poorly understood by the public. I am looking to conduct interviews with representatives of prominent, publicly-traded companies, like yourself, to correct these misunderstandings about the political challenges faced by companies in your industry. As such, I am hoping to capture your perspective on the political landscape that [NAME] faces, which would offer invaluable first-hand information about not only your company, but also about the concerns businesses face more generally.

I expect that the interview will take no longer than 30 minutes of your time. Your responses to the questions will be kept confidential. Neither your name nor your company's name will be released or linked to your comments.

If you are willing to participate, please suggest a day and time that suits you. If you are not able to participate but you know someone at your company or another company who might be interested in participating, I would be grateful if you would connect us.

The interview will be conducted by phone or video chat. If you have any questions please do not hesitate to ask.

Please feel free to reach out at your convenience. I look forward to hearing from you. Thank you for your time and consideration.

CONTROL:

Good morning!

My name is Jane Sumner and I am an Assistant Professor of Political Science at the University of Minnesota, Twin Cities. I am looking to conduct interviews with representatives of prominent, publicly-traded companies for my ongoing book project on companies and politics in the United States.

If you are willing to participate, please suggest a day and time that suits you. If you are not able to participate but you know someone at your company or another company who might be interested in participating, I would be grateful if you would connect us.

The interview will be conducted by phone or video chat. Your responses to the questions will be kept confidential. Neither your name nor your company's name will be released or linked to your comments.

If you have any questions please do not hesitate to ask.

Please feel free to reach out at your convenience. I look forward to hearing from you.
Thank you for your time and consideration.