

PEER ADVICE ON FINANCIAL DECISIONS: A CASE OF THE BLIND LEADING THE BLIND?

Sandro Ambuehl, B. Douglas Bernheim, Fulya Ersoy, and Donna Harris*

Abstract—We investigate the impact of peer interaction on the quality of financial decision making in a laboratory experiment. Face-to-face communication with a randomly assigned peer significantly improves the quality of subsequent private decisions even though simple mimicry would have the opposite effect. We present evidence that the mechanism involves general conceptual learning (because the benefits of communication extend to previously unseen tasks), and that the most effective learning relationships are horizontal rather than vertical (because people with weak skills benefit most when their partners also have weak skills). The benefits of demonstrably effective financial education do not propagate to peers.

I. Introduction

WHEN making financial decisions, people often seek advice from family and friends rather than from experts.¹ This tendency is most prevalent among those with low levels of financial sophistication, due in part to their mistrust of other information sources and avoidance of professional advisors (Collins, 2012; Helman et al., 2013). A large literature shows that these social interactions *affect* personal financial choices, but there is little direct evidence about the extent to which they improve or degrade the quality of decision making (Hastings et al., 2013, also see section II).

Unfortunately, within any group of family and friends, even the most financially knowledgeable individual may have little or no actual expertise. In addition, the financially unsophisticated decision makers who turn to peers are typically ill-equipped to evaluate the advisor's qualifications, or to distinguish between good and bad advice. The practice of relying on peers for financial advice, therefore, raises the concern that people may receive and credit poor advice—in other words, that the blind may (mis)lead the blind.

To complicate matters further, optimal financial decisions depend on preference characteristics that vary from one person to the next, such as patience and risk aversion. The problem of learning from peers becomes more challenging when

preferences enter the mix (see, e.g., Gagnon-Bartsch, 2017, for a formal model). To the extent informal advice simply encourages mimicry, it can be highly suboptimal even when the advisor happens to make the right choice for themselves, because their preferences may differ from those of the advice-seeker. To benefit from peer-to-peer communication, people must be able to either (i) separate principles from preferences and apply the principles based on their own preferences, or (ii) recognize and mimic those with better information and similar preferences.

In this paper, we study the effect of peer advice on the quality of financial decision making in a laboratory experiment with face-to-face interaction. The choices we study require an understanding of some simple financial principles, but they also implicate personal preferences. In contrast to our approach, previous studies that speak to the effects of peer communication on the quality of financial decision making typically employ tasks with dominant alternatives, which remove preferences from the mix, thereby artificially simplifying the problem of social learning (Hvide & Östberg, 2015; Haliassos et al., 2020).

We use a laboratory experiment for three reasons. First, it allows us to measure decision quality cleanly in domains that implicate preferences.² Second, it offers opportunities for introducing controlled variation in conditions and for monitoring communication, both of which facilitate an examination of the mechanisms through which peer effects operate. Third, through the exogenous assignment of treatments and peers, we avoid well-known econometric challenges associated with the measurement of peer effects in observational data (Manski, 1993). While one might worry that people would spend less time on decisions in the laboratory than in the real world, empirical evidence suggests otherwise. A survey of nonfaculty staff at the University of Southern California, for instance, revealed that 58% of respondents spent less than one hour—the duration of many economic experiments—determining their contribution rates and asset allocations for their retirement savings plans (Benartzi & Thaler, 1999).

In our experiment, subjects make decisions concerning investments that accrue compound interest. Optimal choices depend on subjects' idiosyncratic time preferences. We compare subjects in a *Communication* treatment with others in a *Solitary* treatment. In both treatments, subjects start by making decisions about investments in private. Next, those

Received for publication March 15, 2021. Revision accepted for publication July 18, 2022. Editor: Shachar Kariv.

*Ambuehl: University of Zurich; Bernheim: Stanford University; Ersoy: University of Chicago; Harris: University of Oxford.

We are especially grateful to Michalis Drouvelis and the people at the Birmingham Experimental Economics Laboratory for letting us perform our experimental sessions at their facilities. We also thank James Choi, Florian Ederer, Glenn W. Harrison, and the audiences at various seminars and conferences for helpful comments and discussions. David Zuckerman provided excellent research assistance. This research has been approved by the Stanford IRB in protocol 31364. We gratefully acknowledge funding from the Alfred P. Sloan foundation (grant number G-2017-9017) and from the Stanford Department of Economics.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01269.

¹For example, Benartzi and Thaler (1999) surveyed recently hired staff at the University of Southern California about decisions pertaining to their pension plans and found that most of them did not consult anyone other than their family members. See also Bernheim (1998), Lusardi (2003, 2008), van Rooij et al. (2011), Lusardi and Mitchell (2014).

²We assess the quality of decision making using the notion of Deliberative Competence developed by Ambuehl et al. (2022), which accommodates differences in preferences (unlike dominance-based approaches). See section IIIA for details. Bernheim and Taubinsky (2018) explain how this approach relates to and differs from other approaches.

in the *Communication* condition take part in a face-to-face discussion with a randomly assigned peer about similar investments, while those in the *Solitary* condition study the same investments on their own. Subjects then make additional decisions in private. A third treatment, *Indirect Education*, mirrors the *Communication* treatment, except that, before speaking with a randomly matched peer, half of the subjects view an educational video about compound interest that demonstrably improves the quality of their decisions. This treatment allows us to assess whether the benefits of an effective educational intervention propagate through social interactions.

Our analysis consists of four parts. We begin by asking whether peer interaction can improve the quality of individual decision making in a setting where optimal choices vary widely due to preference heterogeneity. We find that communication is, on average, beneficial, especially for those with relatively little financial sophistication, who are more likely to rely on family and friends in practice. To be clear, we do not claim that peer interaction is *always* beneficial. Indeed, the literature appears to offer counterexamples.³ Instead, our study shows that people can benefit from such interaction even under unfavorable circumstances (specifically, when objectives differ).

How do people manage to derive these benefits? The next two parts of our analysis investigate mechanisms. The second part asks whether these mechanisms involve specific learning or general learning. In this context, specific learning refers to the possibility that people emerge from peer interaction with solutions to the particular decision problems they discussed. In contrast, general learning refers to the possibility that, through peer interaction, people acquire a portable conceptual understanding of a class of decision problems that allows them to solve new problems more successfully. To answer this question, we structure our experiment so that subjects make a collection of decisions after discussing a subset of them. We find that the improvement in the quality of decision making due to peer communication is roughly the same for the decision problems subjects discussed, and for the ones they did not discuss. We conclude that the pertinent mechanisms involve general learning rather than specific learning.

The third part of our analysis asks whether the underlying mechanisms involve vertical or horizontal learning relationships between peers. In a perfectly vertical relationship, information flows from the informed to the uninformed: either one person teaches another (a teacher/pupil relation-

ship) or one person serves as an exemplar for another (a role-model/follower relationship). In a perfectly horizontal relationship, peers work together as equals to solve problems or arrive at applicable principles.⁴ According to the vertical learning hypothesis, the better the decision quality of the discussion partner, the larger the beneficial effect of peer communication on a subject's own decision quality. In contrast, according to the horizontal learning hypothesis, communication is more effective between people who appreciate each others' gaps in knowledge, reasons for confusion, and preferred pace; consequently, the benefits from communication may be largest when peers are most similar, even when greater similarity implies lower skill.⁵ Consistent with the horizontal learning hypothesis, we find that people in the bottom half of the decision-skill distribution experience greater improvements when interacting with others in the bottom half than when interacting with others in the top half. Qualitative analysis of subjects' discussions supports this interpretation, in that similarly skilled partners appear to discuss each problem in greater depth.

The fourth part of our analysis asks whether peer-to-peer communication can augment the effects of beneficial financial education interventions that target limited numbers of consumers by propagating their effects through the population. Based on the preceding discussion, two countervailing mechanisms may be at work. On the one hand, consumers who received an *effective* education intervention (treated consumers) acquire new skills, which they may transmit to others. On the other hand, an effective treatment reduces the similarity between treated and untreated individuals, potentially hindering the transmission process. Treated individuals may also have greater difficulty communicating recently acquired conceptual knowledge.

We use an education intervention that demonstrably improves the quality of decision making for treated subjects. Overall, communication with a treated peer is no more beneficial to untreated participants than communication with an untreated peer. However, there is an important qualification. Communication with a treated peer is more beneficial for decision problems the pair discussed, and less beneficial for decision problems the pair did not discuss. This finding suggests that the beneficial effects of communicating with a treated partner are mostly attributable to specific rather than general learning. An analysis of the content of conversations provides corroborating evidence that the intervention reduces the subjects' perceived similarity. It also shows that subjects spend more time discussing a decision heuristic highlighted in the education intervention (the Rule of 72) and less time discussing general conceptual principles (how compound interest works).

³For example, Kaustia and Knüpfer (2012) show that entry into stock market investing is greater following periods in which peers have higher returns, which suggests that peers may encourage return-chasing. Relatedly, Escobar and Pedraza (2019) find that inexperienced students in classrooms undergoing a large-scale financial education initiative overestimate the value of active trading if they are exposed to peers with large returns on single trades. Heimer (2016) shows that peers can amplify the disposition effect because people are unwilling to concede losses in front of others. Our experiment does not encompass distortions resulting from concern for social image because decisions and consequences are private.

⁴A sizable experimental literature finds that when groups make *collective* decisions, they often perform better than individuals (see Kerr & Tindale, 2004 and Charness & Sutter, 2012 for reviews). Whether the benefits translate to private decisions is unclear.

⁵Hoxby and Weingarth (2005), Booij et al. (2016), and Feld and Zölitz (2017) document effects consistent with this mechanism.

The remainder of this paper proceeds as follows. Section II reviews related research. Section III explains the design of our experiment. Section IV describes our data and performs preliminary analyses. Section V addresses each of the four questions featured above and presents our main results. Finally, section VI explores policy implications and concludes.

II. Related literature

This work relates to three strands of literature. The first concerns peer effects in financial decision making.⁶ Notable studies include Duflo and Saez (2003), Hong et al. (2005), Kaustia and Knüpfer (2012), Bursztyn et al. (2014), Beshears et al. (2015), Cai et al. (2015), Frydman (2015), Hvide and Östberg (2015), Heimer (2016), McCartney and Shah (2023), Ko and Pirinsky (2018), Escobar and Pedraza (2019), Ouimet and Tate (2020), and Arrondel et al. (2022).

Our contribution to that literature is fourfold. First, while much of the literature documents directional behavioral effects, we study the *quality* of financial decision making, which we measure in a precise and theoretically rigorous way, within a domain that implicates preferences. Second, we study peer effects in a controlled laboratory setting, thereby identifying causal influences and investigating associated mechanisms in greater detail and under weaker identifying assumptions than existing work.⁷ Third, we distinguish between a peer's baseline competence and recently acquired skills, showing that people have little ability to transmit the latter. Fourth, we provide causal evidence on the role of peers' characteristics.⁸ Our paper is most closely related to Haliassos et al. (2020), which finds that proxy measures of financial decision-making quality are higher among refugees in Sweden when neighbors have economics or business education, but only when the head of the household is educated or male. They infer that the underlying mechanism involves knowledge transfer rather than mere imitation. Our use of a laboratory experiment permits us to deploy a more theoretically robust and quantitatively precise measure of decision-making quality, and to measure causal effects by creating the necessary exogenous variation in conditions.

⁶See Mobius and Rosenblat (2014) for a general review of social learning in economics.

⁷Two studies offer related evidence concerning mechanisms. Bursztyn et al. (2014) examine peer effects among Brazilian investors. They find that these investors both learn from each other and mimic each others' asset holdings. Similarly, Arrondel et al. (2022) study how social interactions affect stock market participation among a representative sample of French investors. Using survey data, they provide evidence of both mechanisms, but especially for the learning channel.

⁸Our laboratory findings have a few observational counterparts. Ouimet and Tate (2020) find that employees who are poorly (highly) informed about employee stock purchase plans for U.S. public firms are most influenced by others who are poorly (highly) informed. Ko and Pirinsky (2018) find that sociability within a county promoted more conservative demand for housing and more stable real estate prices during the 2008 housing bubble, particularly when the number of financially sophisticated residents in an area was high.

The second related strand of literature concerns financial education; see Hastings et al. (2013) and Lusardi and Mitchell (2014) for reviews, and Fernandes et al. (2014); Miller et al. (2015), and Kaiser et al. (2021) for meta analyses. Some have argued for targeting interventions at influencers and relying on social diffusion to leverage the effects of financial education (see, e.g., Haliassos et al., 2020; Ouimet & Tate, 2020). Because the indirect beneficial effects of education in our experiment arise from mimicry rather than from improved conceptual understanding, our experiment calls the effectiveness of many such diffusion strategies into question. Beneficial diffusion may be limited to the transmission of descriptive information and narratives; it appears that conceptual decision strategies are less likely to propagate through social networks.⁹

Third, we contribute to an experimental literature on peer effects in learning. Our paper is most closely related to Kimbrough et al. (2020), which finds that peer-teaching improves learning about Sudoku puzzles, but that ability-tracking has a detrimental effect. In contrast, Boudreau and McCubbins (2010) find that providing subjects taking a mathematics test with polls of their peers' beliefs about the correct answers leads them to perform less well.¹⁰ One important difference from the current study is that tasks involving puzzles and math problems do not implicate preferences.

III. Design

As shown in table 1, all subjects complete three stages of decision making, numbered 0, 1, and 2. Subjects experience two interventions, numbered 1 and 2, interleaved between the three decision stages. The first intervention involves either financial education or a placebo documentary. The second intervention involves either communication or solitary contemplation.

We randomly assign each subject to one of three experimental conditions. The first, *Solitary*, provides a baseline: subjects view the placebo documentary between stages 0 and 1 and engage in solitary contemplation between stages 1 and 2. The second, *Communication*, departs from the *Solitary* condition by randomly pairing subjects and introducing communication between Stages 1 and 2. The communication focuses on specific decision tasks, labeled *Discussed*. The third, *Indirect Education*, is the same as the *Communication* treatment, except that it replaces the placebo documentary with financial education for one member of each pair.

We also assign each subject to one of two roles, distinguished according to whether the *Discussed* tasks appear

⁹Haliassos et al. (2020) conclude that a social multiplier amplifies the effects of education. However, they identify this effect based on whether peers happen to be educated. Consequently, they cannot rule out the possibility that the multiplier effect pertains to characteristics that are correlated with education, rather than with education itself.

¹⁰As we focus on communication in pairs, our experiment differs from the literature on social learning (e.g., Bikhchandani et al., 1992; Banerjee, 1992; Weizsäcker, 2010; Eyster et al., 2018), which considers (mis)learning from others in larger networks.

TABLE 1.—EXPERIMENT DESIGN

Role	A (Receivers)			B (Senders)	
Treatment label	<i>Solitary</i>	<i>Communication</i>	<i>Indirect Education</i>	<i>Communication</i>	<i>Indirect Education</i>
Stage 0		Decision Problems (Tasks: <i>Simple</i> , <i>Anc</i> ₀)		Decision Problems (Tasks: <i>Simple</i> , <i>Anc</i> ₀)	
Intervention 1		↓ Documentary		↓ Documentary	↓ Education
Stage 1		Decision problems (Tasks: <i>Simple</i> , <i>Complex</i> ₁ , <i>Anc</i> ₁)		Decision problems (Tasks: <i>Simple</i> , <i>Discussed</i> , <i>Anc</i> ₁)	
Intervention 2	Contemplation	↓ Communication w Uneducated	↓ Communication w Educated	↓ Communication w Uneducated	
Stage 2		Decision problems (Tasks: <i>Simple</i> , <i>Complex</i> ₂ , <i>Discussed</i>)		Decision problems (Tasks: <i>Simple</i> , <i>Complex</i> ₁ , <i>Complex</i> ₂)	

All subjects proceed through the experiment in the order indicated by the arrows. Table 2 lists the parameters of the task sets *Complex*₁, *Complex*₂, *Discussed*.

in stage 1 or stage 2. We refer to those who performed the *Discussed* tasks in stage 2, after any communication, as *Receivers* (role A), and to those who performed them in stage 1, before any communication, as *Senders* (role B).¹¹ Note, however, that communication can flow in either direction, or not at all. Also notice that, in the *Solitary* treatment, all subjects are in role A. Our analysis focuses on role A subjects.

We evaluate the effects of communication by comparing the change in outcomes from stage 1 to stage 2 for subjects assigned to the *Communication* and *Solitary* conditions. Because we focus on *Receivers* (role A), we can distinguish between effects on choices in tasks that are familiar to their discussion partner (the *Discussed* tasks), and choices in tasks that neither partner has previously seen. We evaluate the effects of education on the educated by comparing the change in outcomes from stage 0 to stage 1 for *Senders* receiving education versus *Senders* not receiving education. We evaluate the indirect effects of education on social contacts by comparing the change in outcomes from stage 1 to stage 2 for *Receivers* (who never receive education) in the *Indirect Education* condition versus those in the *Communication* condition.

In the following subsections, we provide details concerning the tasks, the interventions, and implementation.

A. Tasks and Measures of Decision Quality

Each round of decision making elicits valuations for a future reward. There are three different types of decisions. Decisions in the *simple frame* require the subject to choose between a sure number of tokens V paid the day of the experiment, and x tokens received in t days. Decisions in the *complex frame* are similar, except that the future reward x is framed as a compound-interest accruing investment: “We will invest y tokens in an account with $r\%$ interest per day. Interest is compounded daily. We will pay you the proceeds in t days.” *Ancillary* decisions are similar to complexly

framed decisions, but subjects decide between the investment and a sure number of tokens paid in t days. Accordingly, each ancillary decision has a single correct solution that is independent of the subject’s preferences. In each case, the subject’s valuation is the number of tokens V that leaves her indifferent between the tokens and the future reward. We elicit the indifference point through once-iterated multiple decisions lists that vary the comparison amount, V .¹²

Measures of decision quality. The simply and complexly framed valuations allow us to assess the quality of decision making using the *Deliberative Competence* method of Ambuehl et al. (2022). Intuitively, someone who fully understands the consequences of her decisions should make the same choice in substantively equivalent decision problems even if they are framed differently. The absolute value of the difference in valuations across equivalent simply and complexly framed tasks, therefore, provides a natural measure of decision quality. Ambuehl et al. (2022) show that one can interpret this difference as the maximal welfare loss a subject incurs by having to make a choice in the complex frame rather than in the simple frame, as judged by the preferences she reveals in the simple frame, for which the consequences of choices are transparent (the *normative benchmark*).

A potential reservation concerning conventional methods of Behavioral Welfare Economics is that conclusions may be sensitive to other biases outside the scope of the analysis. In the present experiment, one might worry that biases such as time inconsistency or an excessive suspicion that the experimenter will renege on promised future payments corrupt choices in both frames. If so, then choices in the simple frame do not reveal “true preferences.” Lipsey and Lancaster (1956) reasoned that the right way to handle these types of issues is to analyze all distortions (here, biases affecting valuations) and all remedies for them simultaneously. Because

¹¹We do not explicitly inform subjects of these roles.

¹²The comparison amounts range from 0 to 109 tokens. The first list in each round has a resolution of 10 tokens; the second list in each round has a resolution of 1 token. Appendix A.4 presents screenshots of the decision screens.

of the practical challenges that approach presents, the overwhelming preference of economists, as revealed by research in this area, is to analyze policies one at a time. Indeed, the dominant strategy is to treat the decision making apparatus as flawless except with respect to the bias targeted by the policy one seeks to evaluate. But the data are not in fact generated by an otherwise faultless decision process. Hence, there is no reason to think that solutions derived in this manner are desirable, as Lipsey and Lancaster (1956) pointed out in a related setting with multiple distortions. Nor can one coherently paste together solutions for different biases derived in this way, because each solution implicitly assumes that the other solutions are unnecessary.

To address this difficulty, Ambuehl et al. (2022) introduce *idealized welfare analysis*. This approach requires the analyst to conduct welfare analysis pertaining to an intervention that seeks to address a particular bias under the assumption that other biases *will be* (but have not yet been) addressed through other policies, such as financial incentives, nudges, and commitment devices. When adopting this perspective, the analyst acknowledges that other biases infect the observed data, but does not contemplate using the policy of interest to treat them, because that policy is just one component of a broader effort to diagnose and correct flaws in decision making. Idealized welfare analysis is useful because it provides a conceptually coherent framework for compartmentalizing policy analysis. It permits analysts to arrive at an overall solution to multiple interacting problems by solving them one at a time.

Implementation of idealized welfare analysis might appear challenging because the welfare concept pertains to behavior in the counterfactual scenario for which all other biases are removed, rather than to observed behavior, which other biases may infect. However, Ambuehl et al. (2022) prove that, under reasonably general conditions, their measure of Deliberative Competence (the absolute difference in valuations across the simple and complex frames), which is based on observed choices (without the other biases removed), approximates the idealized welfare effect (with the other biases removed), up to an invariant multiplicative scalar. In this sense, welfare measures based on the Deliberative Competence method are robust with respect to certain types of biases, known and unknown, outside the scope of the analysis that may corrupt choices in the simple frame.

Task parameters. Each complexly framed task specifies three parameters: investment duration, daily interest rate, and principal. As shown in table 2, our investigation involves three sets of these tasks, labelled *Complex₁*, *Complex₂*, and *Discussed*. Each of these sets includes two subsets of three tasks. Each subset involves a single combination of the investment duration and the interest rate, which we present with three different principal amounts in order to increase statistical power. All payments are denominated in tokens, which are worth £0.20 each.

TABLE 2.—TASK PARAMETERS

Task set	<i>Complex₁</i>	<i>Complex₂</i>	<i>Discussed</i>
Investment duration: 72 days			
Daily interest rate	3%	1%	2%
Principal	{3, 7, 11}	{12, 28, 44}	{6, 14, 22}
# of doublings	3	1	2
Investment duration: 48 days			
Daily interest rate	3%	4.5%	1.5%
Principal	{6, 14, 22}	{3, 7, 11}	{12, 28, 44}
# of doublings	2	3	1

The Deliberative Competence method, which we use to measure decision quality, requires us to compare valuations in the complex frame to valuations in the simple frame. Because each investment in table 2 compounds to approximately 24, 56, or 88 tokens over either 48 or 72 days, there are a total of 6 equivalent simply framed prospects. Subjects provide valuations for a similar set of 6 prospects in each stage (0, 1, and 2). To avoid showing subjects the same set of simply framed prospects three times, we vary the future rewards slightly from one stage to the next.¹³ Each set of ancillary tasks, *Anc₀* and *Anc₁*, includes elicitation for three prospects.¹⁴

Order of decision tasks. In addition to depicting the experiment's overall structure (summarized above), table 1 displays the order in which subjects in the various roles and treatments complete the sets of simply framed, complexly framed, and ancillary tasks.

The purpose of Stage 0 is to establish baselines for financial knowledge and valuations of simply framed rewards. Accordingly, all subjects perform the *Anc₀* and *Simple* tasks.

The purpose of Stage 1 is to evaluate financial knowledge, valuations of simply framed tasks, and valuations of complexly framed tasks after the first intervention. Tasks include *Anc₁*, *Simple*, and a set of complexly framed tasks – *Complex₁* for subjects in role *A*, and *Discussed* for subjects in role *B*. We assess the effects of the intervention on financial knowledge by comparing choices in *Anc₀* and *Anc₁* for Senders.

The purpose of Stage 2 is to elicit valuations of simply framed tasks and complexly framed tasks after communication or contemplation. Tasks include *Simple*, *Discussed*, and *Complex₂* for subjects in role *A* and *Simple*, *Complex₁*, and *Complex₂* for subjects in role *B*. Comparisons between assessed Deliberative Competence (the discrepancies between simply and complexly framed choices) in Stages 1 and 2 allow us to determine the extent to which the second

¹³Specifically, future rewards with a 72-day delay are (26, 59, 90), (25, 58, 92), and (24, 57, 91) in Stages 0, 1, and 2, respectively. The corresponding rewards with a 48-day delay are (24, 58, 89), (24, 57, 90), and (25, 58, 91).

¹⁴The parameters of problems in set *Anc₀* in the format (duration, interest rate, principal, future reward) are given by (18, 8, 22, 88), (36, 4, 6, 24), and (54, 4, 7, 56). For set *Anc₁*, they are (18, 4, 12, 24), (36, 6, 7, 56), and (54, 2.67, 22, 88).

intervention affects decision making quality. For subjects in role A, we can separately evaluate the effects on the quality of decision making for tasks they have discussed (*Discussed*) and for tasks they have not previously seen (*Complex₂*).

In all stages, we intermingle decisions in an individually randomized order. Substantively equivalent simply framed and complexly framed decision tasks are never identified as such.

B. Interventions

Communication. Subjects in the nonsolitary conditions physically move next to a randomly assigned partner to engage in private face-to-face communication. We ask them to discuss specific financial decisions. To facilitate these conversations, we distribute six sheets to each pair, each describing one of the complexly framed tasks in the *Discussed* set.¹⁵ Each subject in role B has made decisions concerning the investments on the decision sheets (in Stage 1) before conversing with their partner. Subjects in role A have not seen these tasks before the conversation.¹⁶ We recommend that subjects use 15 minutes for discussion, but they are free to end the discussion whenever they like, and can continue with the experiment once they are done. To help subjects break the ice, we ask them to note two questions they plan to ask their partner, and two pieces of advice they may want to give. We provide no explicit incentives for engaging in discussion, but we remind subjects they will complete 18 additional decision tasks in private, which may include the ones we ask them to discuss; consequently, there is a substantial chance that their payment will be determined by one of those decisions. All decision problems are numbered so that subjects can check whether they have seen a problem before. We unobtrusively record all communication, and subjects are aware that we do so.

Subjects in the *Solitary* treatment receive the same discussion sheets as all other subjects, but are asked to consider the investments in private. Nonetheless, at the outset of the experiment, we inform all subjects, including those in the Solitary treatment, that they will discuss the financial investments with a peer. Unbeknownst to the subjects, in the Solitary treatment that conversation occurs only after all decisions are completed. This design feature serves to hold constant the effort subjects expend on their decisions

in anticipation of having to discuss them with someone else. Accordingly, our design isolates the effects of communication per se, and is unconfounded by the fact that the mere expectation of communication may cause some individuals to pay greater attention.¹⁷

Education intervention. The education intervention we employ in the *Indirect Education* condition consists of videos of narrated slide presentations and practice questions with personalized feedback. The slides and narration are based on the section concerning compound interest from a popular investment guide, *The Elements of Investing: Easy Lessons for Every Investor*, by Malkiel and Ellis (2013). It focuses on the rule of 72, a simple method for approximating the time it takes for an investment to double. The narration is verbatim from the text (with a few minor adjustments), while the slides summarize key points.¹⁸ See appendix A.1 for details.

C. Implementation

All instructions are displayed on screen and explained via an audio recording to minimize experimenter effects.¹⁹ Subjects proceed at their own pace. They begin with a short video recording of one of the authors (Bernheim), vouching that we will pay subjects exactly the amount and at precisely the time we promise them.

Because people typically have access to computational tools when making financial decisions, we provide each subject with a calculator that includes a function for exponentials.²⁰ Subjects know from the outset that, at the end of the experiment, we will select one of their decisions at random, and that the selected decision will determine their payment in Amazon gift cards.²¹ Accordingly, subjects have incentives to choose carefully in every decision.

Additional elements of the experiment include tasks that allow us to assess subjects' comprehension of the mechanics of multiple decision lists, as well as survey questions designed to elicit psychological characteristics, financial literacy, accounts of their decision-making processes, and views about their partners. See appendix A.2 for details.

¹⁵The following is an example. "Decision Task 10. We will invest 6 tokens at 2% for 72 days, compounded daily. You will get whatever is in the account after that time. How many tokens would we have to give you today, so you would be just as happy with receiving those tokens today as with receiving the proceeds in the account in 72 days?" The problem described on every sheet shares this same structure.

¹⁶The purpose of this design feature is to ensure that the advice given by those who have received financial education reflects prior efforts to apply the education to their own financial decisions. In practice, adult financial education generally occurs through the workplace, and companies typically provide it to assist with imminent financial decisions, such as those pertaining to retirement investments. Thus, those who have received this education typically have had a chance to apply it, and presumably speak to others from that experience.

¹⁷In a preliminary version of the *Solitary* treatment, subjects did not anticipate that they would communicate with anyone. The quality of financial decision making was generally lower, presumably because these subjects lacked the social motivation for good decision making that was present in the other treatments. Iyengar and Schotter (2008), Blanes i Vidal and Nossol (2011), and Kuhnen and Tymula (2012) document related attempts to save face.


¹⁸Experiment B of Ambuehl et al. (2022) uses the same intervention. To maximize effectiveness, we tested several versions using a series of pilot experiments on Amazon Mechanical Turk.

¹⁹The instructions are reproduced in appendix C.

²⁰Typically, people also have access to the Internet. However, Lusardi and Mitchell (2011) find that only about 20% of a representative sample make use of these tools for real financial decisions, and the web-based experiment in Ambuehl et al. (2022) finds that an equally small proportion of experimental subjects use these tools for their experimental decisions.

²¹We opted for payment in the form of Amazon gift cards because they simplify the disbursement of delayed payments.

FIGURE 1.—DISTRIBUTIONS OF VALUATIONS AND DELIBERATIVE COMPETENCE



Panels A and B: Distribution of normalized valuations for tasks in the simple frame (panel A) and the complex frame (panel B), censored at 2 to increase resolution. Panel C: Distribution of Deliberative Competence in Stage 1, censored at -1 to increase resolution. In all panels, we display within-subject averages across all relevant decisions for Receivers.

IV. Data Collection and Preliminary Analysis

A. Data Collection

Our subjects are undergraduate students at the University of Birmingham, UK.²² University students comprise an important demographic group whose members are just beginning to make important personal financial decisions. However, many of them may be ill-equipped to do so: a mere 45% of our sample correctly answer three standard financial literacy questions (Lusardi, 2008). Students are also a target demographic group for many financial education interventions.²³

We ran sessions from Fall 2015 to Spring 2016.²⁴ On average, sessions lasted 124 minutes, and subjects earned £26.55, including a £12.5 participation fee. We restrict our analysis to subjects who demonstrated comprehension of the experimental procedures (87.12%, or 460 of 528 subjects).²⁵

²²We chose the University of Birmingham because the subject pool is large and diverse in terms of mathematical skills. We are grateful to Michalis Drouvelis for facilitating our use of the Birmingham Experimental Economics Laboratory.

²³In the United Kingdom, financial literacy education for students aged 11–16 became part of the National Curriculum in September 2014, as part of citizenship requirements (House of Commons Library, 2016). Given their ages and the timing of the study, the subjects in our sample were not affected by this policy.

²⁴See appendix A.3 for details.

²⁵90.78% of Senders and 85.67% of Receivers passed the initial comprehension check. We retain subjects who passed the check even if they

Among those subjects, 99 in the *Communication* treatment, 89 in the *Indirect Education* treatment, and 75 in the *Solitary* treatment participated as Receivers (in Role A). Appendix B.1 presents summary statistics.

B. Temporal Discounting

We first present summary statistics on choices in the simple and complex frames. To compare behavior across rewards of different sizes, we normalize valuations. Specifically, if the future reward associated with decision problem d is given by x , and if $V_{j,d}^f$ denotes individual j 's valuation in decision problem d with framing $f \in \{\text{simple}, \text{complex}\}$, we define subject j 's *normalized valuation* as $\delta_{j,d}^f = \frac{V_{j,d}^f}{x}$.²⁶

Figure 1 shows a histogram of within-subject averages of normalized valuations among Receivers. Panel A displays within-subject averages for decisions in the simple frame, separately for Stage 1 and Stage 2. The vast majority of subjects display discount rates lower than 1. Pooling across

were paired with subjects who did not, since the quality of communication does not depend on understanding the experimental procedures. Despite our precautions, four subjects participated twice. These subjects may have had multiple accounts in the participant management system. We identify these subjects by their email addresses, which they had to supply to receive payment via Amazon gift cards. For these subjects, we retain only the data from their first session.

²⁶Because we elicit valuations using multiple price lists, they are interval-coded. We use interval midpoints.

stages, the average subject values £1 at £0.87 on average if it is paid with a delay of 48 days, and at £0.85 if it is paid with a delay of 72 days (with a population-level standard deviation of £0.25 in both cases).²⁷ These discount factors are comparable to those in the literature on discounting over similarly brief time frames (Frederick et al., 2002). Crucially, the Deliberative Competence approach we employ to measure decision quality yields valid measures of idealized welfare losses even if these choices reflect not only normatively valid discount rates, but also ancillary decision-making imperfections such as present bias or exaggerated concerns that the experimenter will not follow through on the promised delayed payment.

Panel B shows the distribution of valuations in complexly framed tasks, again averaged across all decisions in Stages 1 and 2, respectively. Valuations are significantly more dispersed than in the simple frame, with lower averages of £0.78 and £0.75 for the 48 and 72 day time frames, respectively. The fact that mean valuations are lower in complexly framed tasks than in simply framed tasks is consistent with exponential growth bias, the well-established tendency to underestimate compound interest (Stango & Zinman, 2009). Tellingly, complexly framed valuations in Stage 2 are higher and more tightly concentrated than those in Stage 1, as one would expect if the intervening communication improved decision quality.

C. Deliberative Competence

We measure the quality of subjects' decision making by assessing their *Deliberative Competence* (Ambuehl et al., 2022). Given normalized valuations $\delta_{j,d}^{complex}$ and $\delta_{j,d}^{simple}$, we define Deliberative Competence as $c_{j,d} = -|\delta_{j,d}^{complex} - \delta_{j,d}^{simple}|$. With this sign convention, a higher number indicates greater Deliberative Competence. To interpret the magnitude of $c_{j,d}$, suppose a subject is willing to pay £0.8 for a complexly framed investment that she would value at £1 if she properly understood her opportunity set. In that case, $c_{j,d} = -0.2$. If we are willing to interpret the subjects' choices in the simple frame as normatively valid, we would conclude that she is vulnerable to a maximum welfare loss of £0.20 per dollar due to her deficient command of compound interest. If we are unwilling to accept choices in the simple frame as normatively valid due to suspicion that other biases are present, we can nevertheless adopt an idealized welfare perspective and conclude that, in a setting with effective remedies for other bias, the subject would be vulnerable to a maximal welfare loss equal to 20% of the dollar-equivalent value of the opportunity, as Ambuehl et al. (2022) show.

Panel C of figure 1 shows a histogram of Deliberative Competence averaged over all decisions in Stage 1 for Re-

ceivers. The distribution is skewed, with a mean of -0.26 and a median of -0.18 . The first and third quartiles are -0.42 and -0.04 , respectively. Deliberative Competence is virtually perfect (between 0 and -0.01) for 10% of subjects.

V. Main Analysis

Our analysis proceeds in four parts. We first study the overall effect of communication with a randomly assigned peer on the quality of financial decision making (section VA). We then examine mechanisms. We investigate whether the beneficial effects of communication reflect genuine learning or if they are merely a consequence of subjects' ability to identify high-competence peers and mimic their choices without comprehending the underlying concepts (section VB). We also study how the effectiveness of communication depends on the relative financial sophistication of the subjects and their partners (section VC). Finally, we investigate whether the effects of beneficial financial education propagate through peer-to-peer communication (section VD). In all regressions, we include controls for subject characteristics. Specifically, our controls consist of gender, age, age-squared, ethnicity indicators, an indicator for whether English is the subject's first language, an indicator for whether the subject is an international student, and indicator variables for whether the subject lives in a rural, suburban, or urban area. Appendix sections B.2 and B.3 show that our results are robust with respect to the inclusion of additional control variables, and to the omission of controls.

A. Overall Effect of Communication on Decision-Making Quality

We begin by studying whether communication about financial decisions is helpful, harmful, or haphazard (a case of the blind leading the blind).

Panel A of figure 2 displays the levels of Deliberative Competence across Stages 1 and 2 separately for the *Solitary* and *Communication* treatments. While competence increases slightly in the *Solitary* treatment, the improvement in the *Communication* treatment is substantially larger.²⁸ In our experiment, communication with a randomly selected peer has a clear beneficial effect on the quality of decision making.


To formalize this result, we estimate regressions in which the dependent variable measures the extent to which a subject's Deliberative Competence *changes* between Stage 1 (which precedes communication) and Stage 2 (which follows communication). We consider changes rather than levels to control for individual-level heterogeneity in Deliberative Competence and thus obtain more precise estimates. Specifically, we pair each complexly framed decision d in

²⁷The respective figures are 0.87 and 0.84 for Stage 1, and 0.89 and 0.86 for Stage 2. Analogous statistics for subjects in the role of Sender are similar.


²⁸The increase in Deliberative Competence within the *Solitary* condition does not necessarily reflect learning. An alternative explanation is that Stage 2 employs a different set of decisions than Stage 1.

FIGURE 2.—OVERVIEW OF RESULTS


A. Overall effect of communication.




B. Effect on discussed problems.




C. Effect on new problems.



D. Effect on Receivers in top half



E. Effect on Receivers in bottom half



The panels depict average Deliberative Competence in Stage 1 and Stage 2 for different groups of Receivers, and for different types of questions. Estimates control for demographics; displayed levels pertain to the average subject. See tables 3 and 4 for associated regression estimates.

Stage 2 with the unique complexly framed decision d' in Stage 1 that has the same time frame and (approximately) the same future value, and we let $c_{j,d,1}$ and $c_{j,d',2}$ denote the measures of Deliberative Competence corresponding to these decisions. We then define j 's improvement on task d as $Improvement_{j,d} = c_{j,d,2} - c_{j,d',1}$.

This procedure yields 12 observations per subject, one for each complexly framed decision in Stage 2. We run an OLS regression of this outcome variable on treatment indicators,

using data on Receivers, pooling across discussed decisions (the *Discussed* set) and novel decisions (the *Complex₂* set):

$$Improvement_{j,d} = \beta_0 + \beta_1 Communication_j + \beta_2 IndirectEducation_j + X_j \Gamma + \Phi_d + \epsilon_{j,d} \quad (1)$$

In addition to the subject characteristics listed at the beginning of this section, the vector of subject-specific

TABLE 3.—EFFECT OF COMMUNICATION

Variables	(1)	(2)	(3)	(4)	(5a)	(5b)	(6a)	(6b)
	Improvement in receivers' deliberative competence before/after communication							
Benchmark simply framed choices	Contemp.	Stage 0	Contemp.	Stage 0	Contemporaneous		Stage 0	
Set of decision problems								
<i>Discussed</i>	Yes	Yes	Yes	Yes	No	Yes	No	Yes
<i>New</i>	Yes	Yes	No	No	Yes	No	Yes	No
Counterfactual mimicry	No	No	Yes	Yes	No	No	No	No
<i>Improvement in Solitary</i>	0.021 (0.018)	0.000 (0.019)	0.014 (0.026)	−0.010 (0.027)	0.014 (0.021)	0.027 (0.022)	−0.007 (0.021)	0.007 (0.022)
Row A: Communication								
<i>Improvement</i> <i>(compared to Solitary)</i>	0.066** (0.028)	0.068** (0.028)	−0.107** (0.049)	−0.070 (0.048)	0.065** (0.028)	0.068** (0.029)	0.067** (0.028)	0.069** (0.030)
Row B: Indirect education								
<i>Improvement</i> <i>(compared to Solitary)</i>	0.073*** (0.028)	0.067** (0.028)	0.008 (0.047)	0.033 (0.048)	0.047 (0.029)	0.099*** (0.030)	0.041 (0.029)	0.093*** (0.031)
<i>p-values</i>								
Communication = Indir. Educ.	0.788	0.972	0.039	0.050	0.263	0.519	0.396	0.341
Discussed = Not-discussed if Sender uneducated						0.885		0.918
Discussed = Not-discussed if Sender educated						0.006		0.006
Difference in differences						0.005		0.004
Observations	3,156	3,156	1,572	1,572	526		526	
Subjects	263	263	262	262	263		263	

Improvement in Deliberative Competence from Stage 1 to Stage 2. Based on all subjects in the role of Receiver. Estimates in the *Improvement in Solitary* row indicate the average level of improvement for a Receiver in the *Solitary* condition. Subsequent rows show the additional improvement from communication and indirect education. All regressions control for initial skills and demographic characteristics. See appendix B.2 for additional specifications. The dependent variable in columns (3) and (4) is the hypothetical improvement in Deliberative Competence we would observe if all Receivers blindly mimicked their matched Sender's choices in discussed complexly framed tasks. All other columns use actual improvements. Columns (1)–(4) present OLS regressions with standard errors clustered by subject. Columns (5a) and (5b), as well as (6a) and (6b), each present estimates of a two-equation SUR regression. In the latter regressions, we average improvement within each task set to obtain a single pair of observations per subject. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

controls X_j includes Receivers' preexisting levels of financial skills (measured by their decisions in the Anc_0 and Anc_1 sets, which they made before the communication or contemplation intervention).²⁹ The specifications therefore allow for the possibility that subjects who differ with respect to initial skills may improve to varying degrees over the course of the experiment, either due to differential learning or because of regression to the mean. If we did not control for initial skills, differences in starting points across the treatment groups could manifest as spurious treatment effects. We also include a vector of question fixed effects Φ_d .

Column 1 of table 3 displays the results. The improvement in the *Solitary* treatment, 2.1 percentage points, is statistically insignificant. The increase in the Deliberative Competence of Receivers in the *Communication* treatment is 6.6 percentage points larger than this baseline ($p < 0.05$); see Row A. This size of this effect is particularly impressive when compared to the average level of Deliberative Competence among all Receivers in Stage 1 (26.1%). It bears emphasis that the quantitative interpretation—a communication-induced reduction in the idealized welfare loss of 25.3% (i.e., 6.6/26.1) compared to the initial level—remains valid even if biases other than those pertaining to the evaluation of compound interest afflict choices in both the simple and complex frames (see Ambuehl et al., 2022).

²⁹To be clear, this measure of initial financial skills does not involve the data we use to compute initial Deliberative Competence, which is a component of the dependent variable.

Because Deliberative Competence is defined as the absolute difference in valuations across simple and complex frames, changes in competence could reflect behavioral changes in either frame. To demonstrate that peer communication increases competence primarily through its effect on behavior with complex framing, column 2 replicates column 1 using an alternative normative benchmark to calculate Deliberative Competence: we pair complexly framed choices in each stage with the corresponding simply framed choices made in Stage 0 rather than with those made contemporaneously. By construction, this alternative measure is unaffected by changes in simply framed choices that may result from communication. We find that the estimated treatment effects remain virtually unchanged. Hence, we conclude that communication improves measured Deliberative Competence primarily through its effect on choices in the complex frame.

Significantly, our subjects achieve the gains documented in this section despite the fact that, due to preference heterogeneity, indiscriminate mimicry would be harmful. We simulate mimicry by substituting the Sender's Stage 1 choices for the Receiver's Stage 2 choices in the set of *Discussed* tasks. We then estimate equation (1) using these counterfactual decisions for those tasks. As column (3) shows, indiscriminate mimicry would decrease Receivers' Deliberative Competence by more than 10 percentage points ($p < 0.05$). When we use simply framed choices from Stage 0 as the normative benchmark, statistical significance vanishes, but the coefficient estimate remains negative (column 4).

B. General versus Specific Learning

To understand the mechanisms underlying the beneficial effect of peer communication, we study whether subjects learn something specific about the tasks they discuss, or something they can generalize to new tasks. Panels B and C of figure 2 separate panel A into these two types of tasks (decision sets *Discussed* and *Complex₂*, respectively).³⁰ Comparing discussed and new tasks, the improvement of subjects is essentially the same in the *Communication* treatment, and is only slightly different in the *Solitary* treatment. It follows that communication enhances improvements in decision-making quality to roughly the same degree for both types of tasks. This finding suggests that communication improves decision making through general learning, and not merely through learning that is limited to the discussed tasks.

Formally, we measure these effects by estimating the following two-equation system using Seemingly Unrelated Regression (SUR):

$$\begin{bmatrix} \overline{Improvement}_{j,C} \\ \overline{Improvement}_{j,D} \end{bmatrix} = \begin{bmatrix} \beta_{0C} + \beta_{1C}Communication_j + \beta_{2C}IndirectEducation_j + X_j\Gamma + \epsilon_{j,C} \\ \beta_{0D} + \beta_{1D}Communication_j + \beta_{2D}IndirectEducation_j + X_j\Gamma + \epsilon_{j,D} \end{bmatrix} \quad (2)$$

where $\overline{Improvement}_{j,C}$ is the average improvement of Receiver j across the six *Complex₂* decision tasks, and $\overline{Improvement}_{j,D}$ is the average improvement of subject j across the six *Discussed* decision tasks. We average across the respective decisions to obtain a single pair of observations per subject.

Columns 5a and 5b of table 3 display the results. Row A column 5a shows that a Receiver in the *Communication* treatment improves by an additional 6.5 percentage points for new questions (the *Complex₂* set) compared to the *Solitary* treatment. The corresponding figure for discussed questions (the *Discussed* set) is 6.8 percentage points (row A column 5b). While both of these coefficient estimates are statistically significant at the 5% level, they are not statistically distinguishable from each other ($p > 0.8$). The results remain qualitatively unchanged when we isolate the portion of these effects due to changes in behavior in the complex frame by using the simply framed choices from Stage 0 as the normative benchmark (columns 6a–6b).

C. Horizontal versus Vertical Learning Relationships

The mechanisms underlying the effects of peer communication likely depend on the characteristics of the subjects in each pairing. On the one hand, communication may facilitate the transmission of financial decision-making skills

from those who have them to those who do not. Under this first hypothesis, a subject's improvement should be larger the more competent her peer. On the other hand, effective skill transmission may require an ability to understand each others' sources of confusion and to address each other's questions and concerns at the appropriate level, and at a comprehensible pace. It may also require people to feel comfortable asking questions without fear of embarrassing themselves (Edmondson, 1999), and the act of explaining material to somebody else may improve one's own comprehension. Under this second hypothesis, the benefits from communication may be largest when peers are most similar, even when greater similarity requires the peer to have less Deliberative Competence.

To determine which of these hypotheses more accurately accounts for our data, we classify each subject according to whether her Deliberative Competence falls into the top or bottom half of the distribution before she communicates with another subject. We then study how the effects of communication vary across the four possible combinations of Sender and Receiver types. We perform this classification using the tasks in *Anc₀* and *Anc₁*, rather than the tasks we use to define our primary outcome variables (i.e., the decisions in *Complex₁*, *Complex₂*, *Discussed*, and the corresponding simply framed choices).³¹ Accordingly, our measured changes in Deliberative Competence do not simply reflect regression to the mean.³² To increase statistical power, we pool across the *Communication* and *Indirect Education* treatments.³³

Panels D and E of figure 2 display the results. Panel D shows that there is little room for improvement among Receivers whose initial skills are above the median. Hence, communication has little if any effect on their degree of improvement regardless of the peer's skills. In contrast, Receivers whose initial skills are below the median improve substantially when they communicate, as shown in panel E. Significantly, the improvement is *smaller* for low-skill Receivers paired with high-skill rather than low-skill Senders. Thus, peer-to-peer communication transmits financial decision-making skills most effectively when peers are equally uninformed, rather than when an informed decision maker teaches an uninformed peer.³⁴

³¹For Senders, we also use decision set *Discussed* for classification. Senders make these decisions before meeting their peer, and the inclusion of these decisions increases statistical power.

³²Our analysis is not confounded by regression to the mean as long as measurement error in tasks *Anc₀* and *Anc₁* is uncorrelated with measurement error in our primary outcome variables based on the decisions in *Complex₁*, *Complex₂*, *Discussed*, and the corresponding simply framed choices.

³³Section VD shows that Receivers' choices are similar across these two treatments.

³⁴The following conversation between two of our subjects illustrates the importance of responding to the partner at the right level and pace: *Subject 1: But you already have one whole pie. I hope I'm making it clear. So you've got a whole pie, right? This is day zero. You've got to have a pie, but after day one, you gain a slice of that pie, so you have more slices. And on day two, you get even more slices of pie. Subject 2: Okay, that seems to make sense. Subject 1: Yeah, so that's why you have to add one to it, because you already have this pie. This is one, but this is 0.02. Subject*

³⁰We use the same Stage 1 decisions in both these panels.

TABLE 4.—EFFECT OF COMMUNICATION BY PAIR CHARACTERISTICS

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Improvement in receiver's deliberative competence							
Benchmark choices in simple frame	Contemporaneous				Stage 0			
Sets of decision problems								
<i>Discussed</i>	Yes	No	Yes	Yes	Yes	No	Yes	Yes
<i>Complex</i> ₂	Yes	Yes	No	No	Yes	Yes	No	No
Counterfactual mimicry	No	No	No	Yes	No	No	No	Yes
<i>Improvement in Solitary condition</i>	0.031*	0.023	0.040**	0.042	0.010	0.000	0.020	0.017
<i>for bottom half Receiver</i>	(0.016)	(0.016)	(0.017)	(0.019)	(0.017)	(0.017)	(0.018)	(0.021)
<i>Additional improvement from communication if</i>								
Receiver bottom half								
and Sender bottom half (β_1)	0.164***	0.148***	0.179***	−0.013	0.183***	0.158***	0.208***	0.023
	(0.045)	(0.049)	(0.046)	(0.075)	(0.045)	(0.049)	(0.046)	(0.076)
and Sender top half (β_2)	0.083*	0.054	0.112**	0.184***	0.053	0.029	0.078	0.202***
	(0.044)	(0.044)	(0.049)	(0.058)	(0.044)	(0.044)	(0.050)	(0.056)
Receiver top half								
and Sender bottom half (β_3)	0.005	0.008	0.002	−0.391***	0.005	0.009	0.002	−0.351***
	(0.019)	(0.021)	(0.020)	(0.064)	(0.017)	(0.018)	(0.019)	(0.059)
and Sender top half (β_4)	−0.022	−0.022	−0.021	−0.113***	−0.013	−0.006	−0.019	−0.077***
	(0.020)	(0.022)	(0.021)	(0.027)	(0.018)	(0.020)	(0.018)	(0.028)
<i>p-values about effect on Receiver</i>								
(<i>R</i> bottom, <i>S</i> bottom) = (<i>R</i> bottom, <i>S</i> top)	0.092	0.065	0.181	0.016	0.006	0.010	0.008	0.023
(<i>R</i> bottom, <i>S</i> bottom) = (<i>R</i> top, <i>S</i> bottom)	0.001	0.008	0.000	0.000	0.000	0.004	0.000	0.000
(<i>R</i> bottom, <i>S</i> top) = (<i>R</i> top, <i>S</i> top)	0.029	0.121	0.012	0.000	0.160	0.486	0.057	0.000
(<i>R</i> top, <i>S</i> bottom) = (<i>R</i> top, <i>S</i> top)	0.191	0.157	0.294	0.000	0.321	0.469	0.253	0.000
Joint insignificance	0.004	0.022	0.002	0.000	0.001	0.019	0.000	0.000
Observations	3,156	1,578	1,578	1,572	3,156	1,578	1,578	1,572
Subjects	263	263	263	262	263	263	263	262

Based on all subjects in the role of Receiver. Estimates in the *Improvement in Solitary* row indicate the average level of improvement for a bottom-half Receiver in the *Solitary* condition. Subsequent rows show the additional improvement from communication. The dependent variable in columns 4 and 8 is the hypothetical improvement in Deliberative Competence we would observe if all Receivers blindly mimicked their matched Sender's choices in discussed, complexly framed tasks. All other columns use actual improvements. All regressions control for initial skills, top-half Receiver dummy, decision problem fixed effects, and demographic characteristics. See appendix B.3 for additional specifications. Standard errors are clustered by subject. * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

We formalize these comparisons by estimating the following OLS regression using data on Receivers:

$$\begin{aligned}
 \text{Improvement}_{j,d} = & \beta_0 + \beta_1(R_bottom_j \times S_bottom_j) \\
 & + \beta_2(R_bottom_j \times S_top_j) \\
 & + \beta_3(R_top_j \times S_bottom_j) \\
 & + \beta_4(R_top_j \times S_top_j) \\
 & + \beta_5 R_top_j + X_j \Gamma + \Phi_d + \epsilon_{j,d} \quad (3)
 \end{aligned}$$

where R_bottom_j and R_top_j indicate whether Receiver j is in the bottom or top half of the skill distribution, respectively, and S_bottom_j and S_top_j are defined similarly for the Sender paired with Receiver j . For Receivers in the Solitary treatment, we set $S_bottom_j = S_top_j = 0$. We control for initial skills, decision problem fixed effects, and the demographic characteristics listed at the beginning of this section.³⁵ We cluster standard errors at the subject level.

Results appear in column 1 of table 4. Communication produces the largest improvement in decision quality when a Receiver in the bottom half of the skill distribution interacts with a Sender in the bottom half of the skill distribu-

tion (16.4 percentage points over and above the gains in the *Solitary* condition, $p < 0.01$). While Receivers also benefit from communication with Senders in the top half, the gain is only about half as large (8.3 percentage points, $p < 0.1$). The difference between these estimates is significant at the 10% level. Columns 2 and 3 perform this analysis separately for the sets of new and discussed decision problems, respectively. The results are similar to those in column 1. When Receivers in the bottom half of the skill distribution communicate with similarly skilled Senders, they achieve substantial gains both in the tasks they discussed and in new tasks. In contrast, when those Receivers communicate with more skilled Senders, their gains are considerably smaller, and mostly limited to the tasks they discussed. In contrast, communication does not appear to benefit Receivers in the top half of the skill distribution irrespective of the type of task or their partner's skill. As mentioned previously, this result may reflect the fact that this group's high level of Deliberative Competence limits the potential for gain.

We compare these results to the improvements we would observe if Receivers indiscriminately mimicked their matched Senders. Specifically, we reestimate model equation (3) replacing the Receiver's Stage 2 choices with the Sender's Stage 1 choices for all discussed complexly framed tasks. Column 4 shows that indiscriminate mimicry would yield large gains for low-skill subjects who interact with high-skill subjects, but would not help low-skill subjects who interact with low-skill subjects. (It also shows that

2: On top of that. Subject 1: Yeah. The pie is good example? Subject 2: Yeah, that was much easier. Later in the conversation: Subject 2: You've just taught me more maths than I've ever learned, ever.

³⁵We obtain similar results for specifications with different sets of control variables; see appendix B.3.

TABLE 5.—DISCUSSION CONTENT BY TREATMENT AND PAIR CHARACTERISTICS

Variables	(1) Highlight similarities	(2) Minutes discussed	(3) # problems discussed	(4) # small talk topics	(5) One person proclaims skills	(6) Rule of 72 discussed	(7) Compounding formula discussed
Communication treatment							
Different skills (α_1)	0.341 (0.073)	8.264 (0.799)	4.136 (0.280)	0.659 (0.111)	0.250 (0.068)	0.000 (0.048)	0.744 (0.073)
Similar skills (α_2)	0.622 (0.073)	10.154 (0.775)	4.073 (0.290)	0.400 (0.110)	0.159 (0.068)	0.047 (0.048)	0.698 (0.073)
Indirect Education treatment							
Different skills (α_3)	0.386 (0.073)	8.516 (0.816)	4.023 (0.280)	0.432 (0.111)	0.432 (0.068)	0.773 (0.048)	0.409 (0.072)
Similar skills (α_4)	0.357 (0.075)	9.499 (0.844)	4.476 (0.287)	0.214 (0.114)	0.366 (0.071)	0.786 (0.049)	0.452 (0.074)
<i>p-values</i>							
Effect of similarity							
Comm. tr. ($\alpha_1 = \alpha_2$)	0.007	0.091	0.876	0.099	0.346	0.495	0.652
Indirect Educ. tr. ($\alpha_3 = \alpha_4$)	0.781	0.403	0.260	0.172	0.502	0.849	0.675
Effect of indirect education							
Similar skills ($\alpha_2 = \alpha_4$)	0.012	0.568	0.325	0.241	0.036	0.000	0.019
Dissimilar skills ($\alpha_1 = \alpha_3$)	0.662	0.826	0.775	0.149	0.061	0.000	0.001
All four parameters equal	0.023	0.301	0.674	0.050	0.026	0.000	0.001
Diff-in-diff	0.036	0.576	0.365	0.852	0.856	0.728	0.538
Observations	175	188	171	175	173	172	172

Based on subjects in the *Communication* and *Indirect Education* treatments. Regressions for communication outcomes correspond to equation (4). The Qualtrics survey automatically provides the measure of *Minutes discussed* for all 188 pairs. The other dependent variables are based on classifications of audio transcriptions for 175 pairs. Columns 3, 5, 6, and 7 exclude pairs for which the transcriber encoded the corresponding variable as “unclear.”

indiscriminate mimicry would be harmful for high-skill subjects.) Thus, communication produces the greatest gains precisely when mimicry offers the smallest advantages.

To evaluate the extent to which our results capture the impact of communication on complexly framed choices as opposed to simply framed choices, columns 5–8 replicate columns 1–4 using the simply framed choices from Stage 0 as the normative benchmark. This modification strengthens our conclusions: the incremental improvement when a Receiver in the bottom half communicates with a Sender in the bottom half rather than a Sender in the top half is now statistically significant at the 1% level.

Next, we check whether the foregoing inferences are consistent with the content of subjects’ conversations. To make this content amenable to statistical analysis, we obtained two independent transcriptions of each audio recording using workers hired through Amazon Mechanical Turk. Most recordings were of sufficient quality to allow transcription, yielding qualitative data on the nature of communication for 175 out of 188 pairs in the *Communication* and *Indirect Education* treatments. Research assistants at Stanford University then reviewed these transcripts and coded qualitative information concerning each discussion regarding (i) whether subjects highlight similarities between each other, through statements such as “I’m bad at this too, so let’s see whether we can help each other out,” (ii) the number of decision problems discussed, (iii) the number of prespecified small-talk topics that came up during the discussion,³⁶ (iv) whether one person explicitly claimed to know how to solve the decision tasks, and (v) whether subjects discussed the

rule of 72, which is the focus of our education intervention.³⁷ In addition, each Receiver’s survey interface automatically recorded the time that passed between the beginning of the discussion stage and the resumption of the survey.

For each of these variables, we estimate the following OLS regressions to determine how the composition of a Receiver-Sender pair with respect to Deliberative Competence relates to the types of topics they discussed. Formally, letting Y_i denote a generic communication outcome variable for pair i , we estimate

$$Y_i = \alpha_1 \text{ComDif}_i + \alpha_2 \text{ComSame}_i + \alpha_3 \text{IndEduDif}_i + \alpha_4 \text{IndEduSame}_i + v_i \quad (4)$$

where $\text{ComDif}_i = 1$ if the pair members are in the *Communication* treatment and in different halves of the skill distribution, $\text{ComSame}_i = 1$ if the pair members are in the *Communication* treatment and in the same half of the skill distribution, $\text{IndEduDif}_i = 1$ if the pair members are in the *Indirect Education* treatment and in different halves of the skill distribution, $\text{IndEduSame}_i = 1$ if the pair members are in the *Indirect Education* treatment and in the same half of the skill distribution.³⁸ We use data on subjects in the *Communication* and *Indirect Education* treatments.

The first two rows of table 5 show the estimated coefficients for subjects in the *Communication* treatment.

³⁷Coders also recorded additional information, such as whether the subjects discussed market interest rates. Our research assistants used the summary information form shown in appendix A.5.

³⁸We can match each of our audio recordings to a unique pair of subjects in our data, but we cannot reliably identify whether a speaker participated as a Sender or Receiver. Hence, all our analysis of discussion content is at the pair level.

³⁶Transcribers indicated which of three prespecified small-talk topics subjects discussed: place of origin, field of degree, and years of study.

Column 1 shows that partners were much more likely to highlight similarities when they fell into the same half of the skill distribution than when they fell into different halves (62.2% versus 34.1%, $p < 0.01$). Moreover, column 2 reveals that similar pairs spent roughly 25% more time discussing tasks than heterogeneous pairs (10.15 min. versus 8.26 min., $p < 0.1$ for differences across pairs). This difference relates to the quality and focus of the conversation: while the number of tasks discussed does not differ across similar and heterogeneous pairs (column 3), the number of small-talk topics discussed is smaller when skills are similar (column 4, $p < 0.1$). Appendix B.4 explores the robustness of these conclusions by adding transcriber fixed effects. The estimates and their statistical significance remain largely unaffected.

In the appendix, we extend this analysis in two directions. First, in appendix B.5, we ask whether communication with peers enhances confidence in decision making. Our strategy is to estimate regressions similar to those in table 4, using a subjective measure of confidence as the dependent variable in place of Deliberative Competence.³⁹ We find that the confidence of low-skill subjects increases significantly when they interact with high-skill subjects. The effect is smaller and statistically insignificant when low-skill subjects interact with low-skill subjects.⁴⁰ Thus, communication produces the greatest gains for low-skill subjects when its impact on confidence is smallest. This finding casts doubt on the possibility that communication with peers improves decision making by enhancing confidence.

Second, in appendix B.6, we ask whether a subject's ability to detect financial competence in peers is related to their own competence. At the end of our experiment, subjects answered the following question (on a seven-point scale): "Do you feel your partner had a firm grasp of how to make good decisions in this study?" We regress these responses on a measure of the Sender's actual Deliberative Competence, allowing the effect to differ according to the Receiver's Deliberative Competence. We find that, when partners are more competent, subjects generally recognize them as such. However, we do not detect a robust relationship between the ability to recognize competence in others and the Receiver's own competence. Thus we do not find support for the hypothesis that an inability to recognize competence among those with low competence contributes to their reliance on family and friends rather than those with greater expertise.

D. Indirect Effects of Financial Education

The beneficial effects of peer communication raise the possibility that social networks may propagate the influence of financial education through the population, magnifying its

benefits. Accordingly, in this subsection we examine the indirect effects of financial education on those who have not participated themselves. The preceding results suggest two competing hypotheses. On the one hand, treated individuals acquire new skills, which they may transmit to others. On the other hand, an effective educational treatment reduces the similarity between treated and untreated individuals, potentially stymying the transmission process. Subjects may also have greater difficulty communicating recently acquired conceptual knowledge.

Before investigating the indirect effects of education, we demonstrate that our intervention has a direct beneficial effect on those who participate. We regress Senders' Deliberative Competence in Stage 1 on an indicator for whether they participated in the education intervention. We find that the direct effect of the intervention is to raise Senders' Deliberative Competence. The level of competence is -0.192 for Senders who undergo the intervention, compared with -0.274 for those who do not. The difference of 8.2 percentage points is substantial, and corresponds to roughly a one-third reduction in the idealized welfare loss. The estimated treatment effect is significant at the 10% level. It increases and becomes significant at the 5% level once we control for preexisting skills and demographics; see appendix B.7.⁴¹

Next we test whether peer communication is more effective if the Sender has participated in the education intervention. The pertinent results appear in row B of table 3, which reports estimates based on equation (1). Column 1 shows that there are no detectable indirect effects of education on a Receivers' degree of improvement beyond those arising from communication alone when we pool discussed and new tasks ($p = 0.788$). We obtain similar results when we use the simply framed choices made in Stage 0 as the normative benchmark (column 2).

Differences do emerge, however, once we distinguish between discussed and new tasks. Focusing again on row B of table 3, columns 5a and 5b show that Receivers who are indirectly exposed to the education intervention improve to a significantly greater degree in decision tasks they have discussed than in new ones ($p < 0.05$ in each specification). The difference-in-differences estimates, which compare improvements of Receivers in the *Communication* and *Indirect Education* treatments for the discussed versus new tasks, are statistically significant ($p < 0.05$). Once again, we obtain similar results when we use the simply framed choices made in Stage 0 as the normative benchmark (columns 6a–6b).⁴² These results suggest that education helps people successfully transmit useful information concerning the specific decisions they discuss with others, but it does not help them transmit generalized knowledge.

³⁹At the end of the study, we asked subjects the following question: "Do you feel you had a firm grasp of how to make good decisions in this study?"

⁴⁰For high-skill subjects, there is no evidence that communication meaningfully impacts confidence.

⁴¹Ambuehl et al. (2022) use the same education intervention for a population of workers on Amazon Mechanical Turk and similarly find that it improves Deliberative Competence from -0.256 to -0.186 . The difference of 7.1 percentage points in that study is highly statistically significant.

⁴²Our results are robust with respect to adding and omitting control variables; see appendix B.2.

To what extent does the content of subjects' conversations reflect these differences? Column 1 of table 5 shows that pairs consisting of members in the same half of the initial skill distribution become much less likely to highlight their similarities if the Sender has participated in the education intervention ($p < 0.05$). At the same time, column 5 shows that one of the members becomes more likely to claim they have pertinent skills ($p < 0.05$), and this result also holds for pairs whose members are in different halves of the initial skill distribution ($p < 0.1$). This finding suggests that education may increase the likelihood that Senders rely on strategies akin to "proof by intimidation," which might help Receivers perform better in the problems they discuss while diminishing the benefits of peer communication for new problems.

An additional possibility is that the contents of the education intervention crowd out other methods of decision making. Indeed, column 6 of table 5 shows that, in the *Indirect Education* treatment, more than three-quarters of the pairs discussed the rule of 72 (the main substantive component of the education intervention) whereas virtually no one did so in the *Communication* treatment ($p < 0.01$). Meanwhile, around 70% of pairs discussed the exact compound interest formula, $\text{future value} = \text{present value} \cdot (1 + r)^t$, if the Sender had not participated in the education intervention, while fewer than half did so in the *Indirect Education* treatment (column 7). If the rule of 72 is an inferior substitute for the exact formula (either because the Sender fails to explain it intelligibly, or because the Receiver fails to apply it correctly), then such crowding out could reduce the effectiveness of peer communication.

VI. Conclusion

We have presented an experiment in which communication about financial decisions between randomly paired subjects leads to genuine improvements in the quality of decision making, measured according to the *Deliberative Competence* method of Ambuehl et al. (2022). We have shown that the improvements reflect conceptual learning rather than mimicry of the choices of those who are better informed. The beneficial effects of communication are especially pronounced in interactions between people who are similarly unskilled, and who seem to be more adept at addressing each others' questions and concerns at the appropriate level and pace. Subjecting one member of each pair to an effective financial education intervention, however, provides no benefits beyond those arising from communication alone. The intervention provides subjects with more skills to transmit, but decreases the effectiveness of skill transmission by creating competence differentials and by crowding out more effective forms of communication.

Some researchers argue for targeting interventions at influencers and relying on social diffusion to leverage the effects of financial education. Because the indirect beneficial effects of education in our experiment arise from mimicry

rather than from improved conceptual understanding, our experiment calls the effectiveness of such diffusion strategies into question. For the same reason, our results caution against promoting rules of thumb that are appropriate for particular segments of the population, but that may propagate to other segments for which they are less well-suited.⁴³

A natural extension of our research would involve the study of peer effects in settings where subjects interact with peers of their own choosing. Another extension would examine interaction among larger groups of individuals. In both settings, subjects' abilities to identify those from whom they can benefit most, and to avoid those who set bad examples, would play important roles in determining the effects of social interaction. A third extension would investigate in greater detail the mechanisms that render communication beneficial. We have not attempted to determine whether communication is helpful because subjects learn from the knowledge of their peers, because they learn by teaching their peers (Cohen et al., 1982),⁴⁴ or because interaction invites subjects to think in new ways (Iyengar & Schotter, 2008). We leave these issues for further research.

⁴³For instance, rules of thumb that encourage paying off or abstaining from debt may compound mistakes arising from debt aversion (Martinez-Marquina & Shi, 2024).

⁴⁴A related concept to learning by teaching is *rubber ducking* in computer science. This method involves debugging code by explaining what it does line by line to an inanimate object such as a rubber duck.

REFERENCES

- Ambuehl, S., B. D. Bernheim, and A. Lusardi, "Evaluating Deliberative Competence: A Simple Method with an Application to Financial Choice," *American Economic Review* 112 (2022), 3584–3626. 10.1257/aer.20210290
- Arrondel, L., H. Calvo-Pardo, C. Giannitsarou, and M. Haliassos, "Informative Social Interactions," *Journal of Economic Behavior and Organization* 203 (2022), 246–263.
- Banerjee, A. V., "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics* 107:3 (1992), 797–817. 10.2307/2118364
- Benartzi, S., and R. H. Thaler, "Risk Aversion or Myopia? Choices in Repeated Gambles and Retirement Investments," *Management Science* 45:3 (1999), 364–381. 10.1287/mnsc.45.3.364
- Bernheim, B. D., "Financial Illiteracy, Education, and Retirement Saving" (pp. 38–68), in O. S. Mitchell and S. J. Schieber (eds.), *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement* (Philadelphia, PA: University of Pennsylvania Press, 1998).
- Bernheim, B. D., and D. Taubinsky, "Behavioral Public Economics" (pp. 381–516), in *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 1 (Amsterdam: Elsevier, 2018). 10.1016/bs.hesbe.2018.07.002
- Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman, "The Effect of Providing Peer Information on Retirement Savings Decisions," *The Journal of Finance* 70:3 (2015), 1161–1201. 10.1111/jofi.12258
- Bikhchandani, S., D. Hirshleifer, and I. Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy* 100:5 (1992), 992–1026. 10.1086/261849
- Blanes i Vidal, J., and M. Nossol, "Tournaments without Prizes: Evidence from Personnel Records," *Management Science* 57:10 (2011), 1721–1736. 10.1287/mnsc.1110.1383
- Booij, A. S., E. Leuven, and H. Oosterbeek, "Ability Peer Effects in University: Evidence from a Randomized Experiment," *The Review of Economic Studies* 84:2 (2016), 547–578.

- Boudreau, C., and M. D. McCubbins, "The Blind Leading the Blind: Who Gets Polling Information and Does It Improve Decisions?" *The Journal of Politics* 72:2 (2010), 513–527. 10.1017/S0022381609990946
- Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman, "Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions," *Econometrica* 82:4 (2014), 1273–1301. 10.3982/ECTA11991
- Cai, J., A. De Janvry, and E. Sadoulet, "Social Networks and the Decision to Insure," *American Economic Journal: Applied Economics* 7:2 (2015), 81–108. 10.1257/app.20130442
- Charness, G., and M. Sutter, "Groups Make Better Self-Interested Decisions," *Journal of Economic Perspectives* 26:3 (2012), 157–176. 10.1257/jep.26.3.157
- Cohen, P. A., J. A. Kulik, and C.-L. Kulik, "Educational Outcomes of Tutoring: A Meta-Analysis of Findings," *American Educational Research Journal* 19:2 (1982), 237–248. 10.3102/00028312019002237
- Collins, J. M., "Financial Advice: A Substitute for Financial Literacy?" *Financial Services Review* 21:4 (2012), 307.
- Duflo, E., and E. Saez, "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics* 118:3 (2003), 815–842. 10.1162/00335530360698432
- Edmondson, A., "Psychological Safety and Learning Behavior in Work Teams," *Administrative Science Quarterly* 44:2 (1999), 350–383. 10.2307/2666999
- Escobar, L., and A. Pedraza, "Active Trading and (Poor) Performance," Policy Research working paper 8767. The World Bank (2019).
- Eyster, E., M. Rabin, and G. Weizsäcker, "An Experiment on Social Mislearning," unpublished manuscript (2018).
- Feld, J., and U. Zölitz, "Understanding Peer Effects: On the Nature, Estimation and Channels of Peer Effects," *Journal of Labor Economics* 35:2 (2017), 387–428. 10.1086/689472
- Fernandes, D., J. G. Lynch Jr., and R. G. Netemeyer, "Financial Literacy, Financial Education, and Downstream Financial Behaviors," *Management Science* 60:8 (2014), 1861–1883. 10.1287/mnsc.2013.1849
- Frederick, S., G. Loewenstein, and T. O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature* 40:2 (2002), 351–401. 10.1257/jel.40.2.351
- Frydman, C., "Relative Wealth Concerns in Portfolio Choice: Neural and Behavioral Evidence," working paper (2015).
- Gagnon-Bartsch, T., "Taste Projection in Models of Social Learning," unpublished (2017).
- Haliassos, M., T. Jansson, and Y. Karabulut, "Financial Literacy Externalities," *The Review of Financial Studies* 33:2 (2020), 950–989.
- Hastings, J. S., B. C. Madrian, and W. L. Skimmyhorn, "Financial Literacy, Financial Education, and Economic Outcomes," *Annual Review of Economics*, 5 (2013), 347–373. 10.1146/annurev-economics-082312-125807
- Heimer, R. Z., "Peer Pressure: Social Interaction and the Disposition Effect," *The Review of Financial Studies* 29:11 (2016), 3177–3209. 10.1093/rfs/hhw063
- Helman, R., N. Adams, C. Copeland, and J. VanDerhei, "The 2013 Retirement Confidence Survey: Perceived Savings Needs Outpace Reality for Many," EBRI Issue Brief 384 (2013).
- Hong, H., J. D. Kubik, and J. C. Stein, "Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers," *The Journal of Finance* 60:6 (2005), 2801–2824. 10.1111/j.1540-6261.2005.00817.x
- House of Commons Library, "Financial and Enterprise Education in Schools" (2016), available at <http://researchbriefings.parliament.uk/ResearchBriefing/Summary/SN06156>.
- Hoxby, C. M., and G. Weingarth, "Taking Race out of the Equation: School Reassignment and the Structure of Peer Effects," unpublished manuscript (2005).
- Hvide, H. K., and P. Östberg, "Social Interaction at Work," *Journal of Financial Economics* 117:3 (2015), 628–652. 10.1016/j.jfineco.2015.06.004
- Iyengar, R., and A. Schotter, "Learning under Supervision: An Experimental Study," *Experimental Economics* 11 (2008), 154–173. 10.1007/s10683-007-9164-2
- Kaiser, T., A. Lusardi, L. Menkhoff, and C. Urban, "Financial Education Affects Financial Knowledge and Downstream Behaviors," *Journal of Financial Economics* 145 (2021), 255–272. 10.1016/j.jfineco.2021.09.022
- Kaustia, M., and S. Knüpfer, "Peer Performance and Stock Market Entry," *Journal of Financial Economics* 104:2 (2012), 321–338. 10.1016/j.jfineco.2011.01.010
- Kerr, N. L., and R. S. Tindale, "Group Performance and Decision Making," *Annual Review of Psychology* 55 (2004), 623–655. 10.1146/annurev.psych.55.090902.142009
- Kimbrough, E. O., A. D. McGee, and H. Shigeoka, "How Do Peers Impact Learning? An Experimental Investigation of Peer-to-Peer Teaching and Ability Tracking," *Journal of Human Resources* 57 (2020), 304–339. 10.3368/jhr.57.1.0918-9770R2
- Ko, K. J., and C. Pirinsky, "Did Social Interactions Fuel or Suppress the US Housing Bubble?" unpublished manuscript (2018).
- Kuhnen, C. M., and A. Tymula, "Feedback, Self-Esteem, and Performance in Organizations," *Management Science* 58:1 (2012), 94–113. 10.1287/mnsc.1110.1379
- Lipsey, R. G., and K. Lancaster, "The General Theory of Second Best," *The Review of Economic Studies* 24:1 (1956), 11–32. 10.2307/2296233
- Lusardi, A., "Planning and Saving for Retirement," unpublished manuscript (2003).
- "Household Saving Behavior: The Role of Financial Literacy, Information, and Financial Education," NBER working paper 13824 (2008).
- Lusardi, A., and O. Mitchell, "Financial Literacy and Planning: Implications for Retirement Well-Being" (pp. 17–39). In A. Lusardi and O. S. Mitchell (eds.), *Financial Literacy: Implications for Retirement Security and the Financial Marketplace* (Oxford: Oxford University Press, 2011).
- "The Economic Importance of Financial Literacy: Theory and Evidence," *Journal of Economic Literature* 52:1 (2014), 1–44.
- Malkiel, B. G., and C. D. Ellis, *The Elements of Investing: Easy Lessons for Every Investor* (Hoboken, NJ: Wiley, 2013).
- Manski, C. F., "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies* 60:3 (1993), 531–542. 10.2307/2298123
- Martinez-Marquina, A., and M. Shi, "The Opportunity Cost of Debt Aversion," *American Economic Review* 114:4 (2024), 1140–1172.
- McCartney, W. B., and A. Shah, "'I'll Have What She's Having': Neighborhood Social Interactions Lead to Policy Spillovers," *Journal of the Association for Consumer Research* 8:4 (2023).
- Miller, M., J. Reichelstein, C. Salas, and B. Zia, "Can You Help Someone Become Financially Capable? A Meta-Analysis of the Literature," *The World Bank Research Observer* 30:2 (2015), 220–246. 10.1093/wbro/lkv009
- Mobius, M., and T. Rosenblat, "Social Learning in Economics," *Annual Review of Economics* 6:1 (2014), 827–847. 10.1146/annurev-economics-120213-012609
- Ouimet, P., and G. Tate, "Learning from Coworkers: Peer Effects on Individual Investment Decisions," *The Journal of Finance* 75:1 (2020), 133–172. 10.1111/jofi.12830
- Stango, V., and J. Zinman, "Exponential Growth Bias and Household Finance," *The Journal of Finance* 64:6 (2009), 2807–2849. 10.1111/j.1540-6261.2009.01518.x
- van Rooij, M., A. Lusardi, and R. Alessie, "Financial Literacy and Stock Market Participation," *Journal of Financial Economics* 101:2 (2011), 449–472. 10.1016/j.jfineco.2011.03.006
- Weizsäcker, G., "Do We Follow Others When We Should? A Simple Test of Rational Expectations," *American Economic Review* 100:5 (2010), 2340–2360.