Pay Attention: Object Consideration as a Mechanism of Network Diffusion

Anthony Vashevko
Stanford Graduate School of Business
vashevko@stanford.edu

April 20, 2015

Abstract

Theories of diffusion in networks rely on two broad classes of mechanisms: social influence directs the flow of information, and influence affects object valuation. This paper proposes an integrated model of decision-making for the adoption process. This model reveals a neglected middle step between information and valuation: social influence affects whether agents consider a particular object as relevant to the adoption decision. I identify this attention-driving mechanism using data on traders in an online foreign exchange platform. Features of the setting exclude the possibility of information- or valuation-driven diffusion, but traders still adopt one another's trading behaviors, and they do so most for those rare behaviors that lack external drivers of attention. I discuss the importance of attention-driven diffusion for future work and the value of an integrated decision model in delineating conditions under which existing theories of influence apply.

Introduction

What moves through the network when an object or practice diffuses? Since the introduction of network analysis into social science, researchers have recognized that the structure of relationships around a person exerts substantial pressure on their choices and outcomes. To date, however, network theorizing has circled around two classes of mechanisms: those of informational diffusion, and those of functional benefit. Networks can convey information about the existence of possible decisions (e.g. learning about new jobs, Granovetter 1973), or networks can make the outcome of a particular decision better or worse, either directly (e.g. email becomes more valuable as more people can be emailed) or by revealing the outcomes of possible decisions (e.g. the quality of partners, Podolny 2001). Yet empirical studies of network diffusion repeatedly hint at gaps in the machinery. Networks can matter even if there is no information to diffuse or valuation to shift.

Where theory falters is by looking at diffusion through disjoint lenses, as a set of disjoint processes. This paper examines diffusion through a model of a decision process – an object or practice diffuses when the adopter selects it as a solution to a problem. Mechanisms of information flow and functional benefit slot in as steps in the decision process. But the model reveals that mechanisms of informational and functional diffusion exclude an important middle step, an attentional account of social influence – influence can work by changing an agent’s consideration of an object, taking options that are well-known and well-understood and pulling them into or out of the agent’s decision set. Between the process of learning about the existence of an object and evaluating the merits of adopting it, the individual must recognize what use the object can be put to or whether the object applies to the decision of the particular moment.

Attention holds the gate to evaluation. In an environment rich with options whose value may be idiosyncratic to the adopter or whose evaluation is expensive, attention mitigates the risks of adoption. In such environments, social influence need not derive from awareness of which options are available or what the value of those options is – instead it allows the decision-maker to learn whether an option is worth learning about, whether they should look for further information (even further information from their network) to make a decision. Though the stylized patterns of diffusion (cascades of adoption, clusters of use) may not differ when attention is the binding constraint, attention separates the use of an object from behaviors of learning about it or consideration of it as an option. When attention binds, researchers may find some parts of a network amenable to adoption and others closed to it, insofar as attention has diffused before use. And they may find clusters of interest and learning remote from clusters of use.

The aim of this paper is to highlight this attentional role of social influence and the value of a coherent adoption model. Social influence is particularly important in helping individuals understand and consider objects they may have known of but never paid attention to. In between the processes of disseminating initial information about some object and producing the structural conditions necessary for its adoption, social networks help individuals make sense of it and enable its use.

This study exploits a dataset on individual currency trading behavior, in which diffusion should not occur but does. As this paper will argue, the setting excludes the possibility of informational diffusion and network effects of adoption, isolating consideration as the mechanism behind object adoption. By demonstrating the ability of social influence to enable the consideration of known objects, this paper aims to establish the existence of attention-influence processes, to show the value of the

adoption model, and to inspire future work on a unified approach to diffusion. Though consideration may not play the driving role in most observed instances of diffusion, pulling apart the mechanisms involved in the adoption process can enable researchers to better understand and predict the causes and courses of such processes.

The article proceeds in four sections: the first discusses existing research on social influence, presents the adoption model, lays out an identification strategy of the model, and discusses the features and merits of the setting. The second describes the empirical strategy of the paper. The third presents the major results of the study, and the fourth concludes.

Mechanisms of Social Influence

Diffusion is the focus of a large body of research (see Rogers 2003 for an overview), and any attempt to discuss it must paint with broad strokes. Studies of object and practice diffusion typically focus on one of two major families of mechanism: the first is exposure to a new object or practice; the second is the existence of structural factors that change valuation during an adoption decision. The first family focuses on pure models of how the structure of a network governs information flow. It goes back to the earliest work on job offers in social networks (Granovetter 1973) and is exemplified by modern work seeking to tease apart the importance of individual actors in information diffusion (Aral and Walker 2012; Muchnik, Aral, and Taylor 2013), or looking to model the flow of news directly (Centola and Macy 2007; Banerjee et al. 2013). The second family focuses on how social influence and network structure affects the decision-making calculus of agents, whether through network effects or uncertainty reduction (Podolny 2001). This style has a natural home in economic study of network games (Jackson 2010), but it likewise shows up in neo-institutional

accounts of practice diffusion (Edelman, Uggen, and Erlanger 1999). These analyses leave something on the table.

Empirical studies of diffusion occasionally run into the problem of describing settings in which the content of possible information flow or network effects is unclear. One example of this problem appears in Davis and Greve’s (1997) study of the corporate diffusion of golden parachutes and poison pills in response to a wave of hostile takeovers. The authors describe a diffusion process driven by the need to discern or justify the legitimacy of diffusing practices: golden parachutes spread through regional networks as boards or managers learned to justify them from local peers; poison pills spread through board interlocks as boards learned about the practice from experienced directors. Yet this explanation remains incomplete. Is there information diffusion at work, and if so, what is the information being spread? Is it knowledge of ways to justify particular practices to boards? Or is there instead a process of re-evaluation: are board members learning that adoption of the practices will be sanctioned by their local or national peers? In general, it is hard to tell who the decision makers are that adopt the practice and how the network affects them at the moment of decision.

To understand the role of social influence in an adoption decision, then, it is helpful to break down the stages of the decision-making process. In particular, it is helpful to consider a model of a nested decision-making process that selectively prunes potential choices until a decision can be made (Roberts and Lattin 1991). An individual faces a decision that can be resolved through some choice (see Figure 1). All possible choices start out in a universal set. First, for a specific individual, this set narrows into an awareness set – those candidates that an individual knows exist. Second, the set narrows to a consideration set – the set of choices that an individual considers relevant to a particular decision. Finally, a single choice is made from within the consideration

set.¹

Transitions between these sets can occur for many reasons, but in particular, they may occur due to social influence. I assume a broad definition of influence: influence need not be intended, and agents can be influenced by peers through learning or observation. Influence can affect each transition. It can affect how objects enter the awareness set, as peers learn from each other the existence of novel options. It can affect how objects enter consideration. Finally, it can affect the evaluation of objects within the consideration set. The first step has been well-explored in studies and theories of information flow. The last step has likewise been covered, typically under one of two explanations: (1) networks directly affect the value an individual can expect from using an object, or (2) information flow can reduce uncertainty around a decision, affecting the expected value of the outcome. But the middle step – the possibility that influence affects which of an agent’s known options enter consideration – has been relatively neglected.

How does influence affect whether an object is considered or not? As Zuckerman (1999) argues, the construction of a consideration set relies on a categorization process – decisions call up categories of solutions, and an object falls into the consideration set for a particular decision if the decision-maker perceives it to be a member of the relevant category. An object’s perceived attributes determine which categories it belongs to. A wine critic recommending top wines may not have considered including California wines until the “Judgment of Paris” revealed that California wines had attributes that made them competitive with French wines, and deserved membership in the class of top wines. Because an object’s attributes are independent of the

¹Roberts and Lattin also discuss the possibility of an intermediate “choice set” that exists just prior to the decision. For the purpose of this article, the distinction between the consideration and choice set is moot.

Figure 1: Steps of the decision process

decisions for which those attributes matter, we can best analyze the effect of influence on the consideration set for a particular decision by looking at how influence affects an object’s attributes.²

Combining these models, I argue that social influence can work to assign and activate categories to objects and thus to place objects into consideration sets. A key mechanism of social influence is the ability of the influenced individual to learn that a particular object has attributes and belongs to categories that render it relevant to a particular decision. Social influence allows individuals to consider or neglect an object in the decisions they make.³

²This mechanism covers even relatively non-“sociological” models of consideration. For example, a key attribute of an object may be the degree of recent exposure to it. Agents may consider objects for no more meaningful reason than that it is particularly salient to them at the moment of decision.

³This idea bears a similarity to March’s logic of appropriateness (1994), whereby individuals behave in the manner that is appropriate to someone like them in the situation that they find themselves in. The consideration process likewise defines the scope of possibilities, though in this case the set of appropriate actions is broader than that of March: an individual may make decisions under a proximately economic logic, but only among those choices that can be considered appropriate.

Returning to the work of Davis and Greve (1997) can help clarify both the steps of this model and the implications of the prior work. The decision-maker is either a board of directors or a manager pitching a potential practice to the board. First, local networks can change the information set in a decision by informing the board of the existence of golden parachutes or poison pills, or by informing the manager of arguments that can be made to legitimize the practices. Arguments made by Davis and Greve suggest that the information flow mechanism is not at work – there was much public discussion of both pills and parachutes. Second, local networks can change the consideration set by informing the board that pills and parachutes are appropriate options for a company of their kind. Finally, local networks can affect the value of a practice by letting a board learn the consequences of adoption from other businesses or learn that their business partners will not punish the company for its adoption of the practice. Davis and Greve do not distinguish between these last two possibilities. It is unclear whether business networks provided information that allowed boards to make better decisions about the adoption of a practice or simply whether local networks changed a board’s perception that the practice was worth its attention.

It is important to point out that this model bears a strong resemblance to March’s analyses of organizational decision-making (e.g. March and Olsen 1976). When evaluation of options is difficult or expensive, constraints on the consideration set can improve an agent’s or an organization’s performance. An alternative view of this article’s core argument is that networks and peer influence provide a key source of these constraints.

Finally, I make no argument that local social influence is the sole driver of object consideration and categorization. Recent work in category theory describes a global

Appropriateness may produce a consideration set consisting of multiple viable options.

category assignment process carried out by some audience (Hannan, Pólos, and Carroll 2007). Similarly, research on sensemaking often describes a global coordination processes as individuals make sense of objects or actions by responding to the actions of all others in a particular field (Antonio et al. 1999). Net of the global exposure to a particular categorization scheme, local social influence may be particularly effective at driving consideration to the extent that peer behavior is particularly salient or offers a particularly high bandwidth of information transmission (cf. Aral and Van Alstyne 2011).

In sum, this paper aims to establish the existence of consideration as a mechanism of social influence and diffusion processes. In the remainder, I lay out the criteria for establishing such a result and produce evidence of a consideration process in a real world setting.

Isolation of Influence Mechanisms

Identifying influence mechanisms in the wild can be difficult because available data typically covers only two aspects of the decision making process: the social structure leading up to a decision and the outcome of the decision itself. Identifying a mechanism thus requires the isolation of intermediate pathways. As discussed above, the adoption decision can be broken down into a series of nested steps. First, an individual learns of the existence of an object. Second, the individual brings that object into their consideration set. Finally, the individual evaluates choices within the consideration set and selects one (or more) to use.

In order to establish the existence of an attentional effect of social influence, analysis must exclude or control for the possibility of social influence at the first or third stages of the decision-making process. That is to say, in observing that the actions of alter

influence ego to adopt some object, the analyst must account for the possibility that alter's actions caused ego to learn of the existence of the object, and the analyst must account for the possibility that alter's actions changed ego's expected outcome in adopting the object. If the research can exclude these two possibilities, the remaining peer effect can be attributed to alters' ability to drive ego's attention toward or away from the object. Achieving these exclusions is the goal of this paper. Notably, studies of other stages of the adoption process rely on similar logic: studies of pure information diffusion must assume that the relevance and value of an object are self-evident, such that adoption follows learning, while studies of the structural dependence of evaluation must assume that all options are in the consideration set, such that adoption proceeds from rational choice.

Identifying Consideration in a Trading Network

In this study, a dataset describing the behavior of individual currency traders on a social trading website imposes key restrictions on the adoption process. On the website, individual traders participate in the global foreign exchange market, placing money on predictions that the price of a particular currency pair will rise or fall over time. At the same time, they can make friends with other currency traders on the platform and observe one another's behavior, allowing for the possibility of social influence.

In line with the discussion above, the currency trading setting allows me to argue that influence on the network does not work by informing traders of the existence of currency pairs to trade in, and it does not work by informing traders of the payoffs to trading in a particular pair or by changing the consequences of trading. Any observed influence effect works by changing a trader's consideration of particular currency pairs.

The following sections discuss the market in greater detail, but briefly I argue for three assumptions: First, the setting excludes the possibility that traders learn about the existence of an object because the set of world currencies is essentially fixed and all possible trading pairs are known up front – there is no network diffusion of object awareness. Second, the foreign exchange market is efficient, precluding the possibility that certain strategies have true value that can be learned from the network – social influence does not affect the expected value of a trade. Third, participation in the market is financially risky and behavior in the market lacks social payoff, so that traders do not trade in order to influence their social relationships – there are no network formation benefits to object use.

This leaves two possibilities for learning, one at the evaluation stage, and one at the consideration stage: it may be that traders engage in superstitious learning based on the stochastic outcomes of their peers, seeing value where none exists; or, as this paper argues, traders may learn about the relevance or appropriateness of particular currency pairs to a good trader’s repertoire – that is, they learn that the currency pairs their peers trade in should be members of their own consideration set when deciding what to trade in. The implicit model is that traders conduct ongoing analysis on pairs in their repertoire and enter the market when analysis indicates that trading would be favorable. Interviews with participants of the platform as well as posts on the online forum support this model of behavior. I consider a trader’s first trade in a particular currency pair to be a meaningful proxy of the trader’s incorporation of that pair into his repertoire – this is the adoption behavior of interest here. As such, I make three predictions about the setting:

First, I expect that observing peer trading in a particular currency pair conveys information about the appropriateness of that currency pair to a trader’s repertoire,

pushing the currency pair into the trader’s consideration set.

Hypothesis 1 Observing any peer trading in a particular currency pair will make a focal trader more likely to adopt trading in the pair.

At the same time, when the process of social influence is transmitting information about currency pair appropriateness and not currency pair value, I expect traders to respond more to the existence of peer trading than to the profitability of their friends’ trades.

Hypothesis 2 Observing peer trading in a particular currency pair will make a focal trader more likely to adopt trading in the pair than will observing peer profits in the currency pair.

Finally, we can exploit one further feature of the setting: Both the distribution of trading in and information about different currency pairs is particularly skewed towards a few popular pairs. Certain pairs are commonly traded and discussed (e.g. the Euro/US dollar trade), while others are relatively rare (e.g. the Euro/Hungarian Forint trade). Because such external information sources act as a source of influence on adoption decisions, and because they strongly favor coverage of popular currency pairs, I expect that traders enter the setting with large gaps in attention to different currency pairs. This setting only allows me to detect the adoption, and not the discontinuance, of currency pairs. As such, the differential degree of externally-driven attention to particular pairs leaves more room for local social influence to operate on uncommon pairs. As such, we expect the social influence to be more effective at increasing consideration of rare pairs than common pairs.

Hypothesis 3 Observing peer trading in a particular currency pair will make a trader more likely to adopt trading in the pair for rarer currency pairs than for common pairs.

Social Trading and Foreign Exchange

The setting for this study is a social trading platform (which I will refer to as SocialTrade). A social trading platform is an online service that allows individual foreign exchange traders to link up their existing trading accounts and conduct trades through the service’s website. The added feature of social trading websites is the ability to send out friend requests to other traders. When two traders agree to become friends, they can observe each other’s trading activity: each can see the activity of the other on his dashboard, the screen that greets him as he opens the site. Figure 2 shows an anonymized mockup of a typical dashboard. Traders can observe the types and outcomes of one another’s trades, though they cannot observe the volume of money being traded.

Social trading websites advertise themselves both as learning and commitment mechanisms: traders can learn from one another’s practices, and traders can motivate themselves to trade real money by having to do it in public.

One top site social trading website, eToro, describes the benefit of social trading as follows:

Thanks to our social investment network, eToro gives you the opportunity to channel the collective potential of a diverse trading community into individual benefit, by learning from, interacting with, and even automatically

copying other network members in real time.

currency	user	position	open time	SL	enter	TP
GBP/USD	Alice	LONG	12:28 AM	NA	1.66061	1.66470
EUR/USD	Bob	SHORT	Jan 22	1.37721	1.35579	1.34964
EUR/JPY	Carol	LONG	Jan 21	NA	141.701	NA
GBP/USD	Alice	SHORT	Jan 21	NA	1.65075	NA
EUR/USD	Alice	SHORT	Jan 21	1.37388	1.36204	1.34388
EUR/USD	Bob	LONG	Jan 18	1.34712	1.35748	1.38112

Figure 2: Example trader dashboard

The reality is bleak. The platform is home to several thousand active traders. These traders appear to engage primarily in short-term momentum trading. It does not appear that traders use the website as a hedging mechanism for a broader portfolio – posts on the site’s forum suggest the same. The median trade is open for 1.5 hours and 90% of trades close within 1.5 days. The mean trader had 8.3 friends and traded in 8.1 currency pairs in his time on the site. Most traders lose money. The mean trader lost about \$17,000 in his time on the site; the median trader lost \$116. Over the 6 years of data on the platform, spanning January 1, 2008 to January 1, 2014, traders engaged in over 4 million trades in 133 currency pairs. Consistent with the broader market, 90% of trading volume occurred in the top 10 currencies, with the most trading occurring in the EUR/USD pair. Nevertheless, 72 currency pairs saw at least 100 trades, and 42 saw over 1000 – there was diversity in traders’ adoption of currency pairs.

As described above, the platform offers a good setting to track individual adoption behavior. Currency pairs are well understood by traders so that the set of all possible tradeable pairs is in some sense known to all traders – it falls into their awareness set.

At the same time, there are no economic network externalities here, nor any possibility of market control. Traders operate in a large market dominated by institutional investors. Though substantial leverage is possible, it is nevertheless unlikely that traders can exert any form of market power, alone or in coalition. At best, traders may adopt currency pairs in order to posture to others, such that currencies serve as signals of some group membership or status. Such posturing appears unlikely given the ever-present threat of substantial monetary losses, and the lack of any apparent benefits to posturing.

The last possible concerns are the possibility of inherent differences in quality among currency pairs or traders. If certain traders and certain currencies systematically outperform others, traders may be smart to follow these superstars and adopt their behaviors. Since such skill is unobservable to the researcher, this study relies on existing research on individual trading behavior, as well as a coarse test for skill, to establish that such differences are unlikely to exist, or, more specifically, are unlikely to be more readily apparent to participants than to researchers.

Structure of the Foreign Exchange Market

The foreign exchange market is a huge and liquid marketplace: daily trading volume exceeds 5 trillion dollars a day, with much of that concentrated in a few key currencies (International Settlements 2013). Prices move quickly, and large institutional investors dominate trading. Insofar as reliable profit requires the exploitation of some rent-generating asset, be that speed, access to information, or sheer market power, individual traders appear to have very little hope of cutting it.

Research on foreign exchange markets is inconclusive on whether profitable strategies or traders exist. As discussed above, most investors on the platform practice short-

term technical analysis (trading based on analysis of price history). Despite possibly violating efficient market theories, technical trading is a popular and persistent strategy in currency markets (Menkhoff and Taylor 2007). A substantial literature has sought to figure out whether technical trading can be profitable. While a number of studies find returns to technical trading rules, this profitability often disappears after including transaction costs (Bessembinder and Chan 1998; C.-H. Park and Irwin 2007; C.-h. Park and Irwin 2010). Moreover, the profitability of successful rules fades rapidly (Schulmeister 2009; Menkhoff et al. 2012). In another vein, studies of individual investors have found that the vast majority are not profitable, especially after controlling for transaction costs and fundamental factors, and those that are profitable retain little of their skill over time (Barber and Odean 2000; Barber et al. 2014; Abbey and Doukas 2014).

It is important to point out, however, that research does suggest that social trading websites can separate out skilled traders from unskilled traders. Heimer and Simon (2012) and Simon (2013) show that traders preferentially make friends with alters that make the best returns, and Pan, Altshuler, and Pentland (2012) show that trades made by the most popular traders outperform those by the least popular. Nevertheless, even the best trades they identify appear to be barely profitable. These results indicate that traders are distinguished mostly by the extent to which they can avoid trading mistakes.⁴ Human beings are prone to a number of cognitive biases that can make them particularly bad risk-takers. In particular, prospect theory (e.g. Kahneman and Frederick 2005) suggests that individuals have inconsistent risk preferences across losses and gains. As one trader suggests in an interview, bad traders tend to take their wins early, and keep their losses open long in hope of a turnaround: the first thing a trader learns is discipline to some program that can prevent losses from accumulating.

⁴See Appendix A for supplementary analysis confirming this story.

While traders believe in the viability of their own strategies, there is little evidence that they are detecting any real or long-term effects.

As such, I argue that the setting minimizes the possibility that following and imitating any particular trader is the basis of a money-making strategy. Though traders may believe they are learning something valuable, they are no more informed than the researcher in this regard. This allows us to argue that any learning that does take place is either superstitious, or works at a different level of the decision-making process: by informing traders that particular currency pairs belong in their consideration set.

Methods

The goal of this study is to identify diffusion through the mechanism of object consideration. This study aims to identify this mechanism by excluding the possibility of learning about objects and learning about profits, and by observing distinct diffusion patterns between common and rare currency pairs: for common currency pairs, their appropriateness to a trader’s repertoire has already diffused; for rare currency pairs, the process is ongoing and mediated by relationships among traders on the platform. Specifically, I test how the influence of peers affects the adoption behaviors of individual traders, and how this affect varies over the set of currency pairs.

Method of Analysis

To test how social influence affects adoption, this study looks at each trader’s decision to adopt a given currency pair on a given day. Adoption here is defined as the trader’s first use of a particular currency pair. In addition, this measure of adoption excludes currency pairs observed on the trader’s first day of trading – if the first trade Alice

made on the platform was in EUR/USD, it is not considered an adoption event, and Alice is excluded from the risk set of EUR/USD adopters.

The unit of analysis is the trader-pair-day: for each trader, each currency pair and each day the trader is on the site, I construct a binary variable indicating whether the trader adopted the given currency on that day. Once a trader adopts a currency, his remaining time on the site, for that currency pair, is censored out. This is similar both conceptually and statistically to an event history, or survival, model.

For computational tractability, I elect to estimate this process in a logistic regression framework instead of a more typical hazard model. There are over 500 million trader-pair-days observed. The introduction of daily covariates in the primary and robustness analyses make splitting and managing the spells inconvenient for a conventional Cox or accelerated failure time model. Instead, I estimate a model equivalent to a Cox proportional hazards model in a conditional logistic regression framework.

The proportional hazards model assumes a time-varying baseline hazard rate for the occurrence of some event, which is then multiplied in the presence of covariates. With discrete time intervals and multiple events per day, one way to implement this model is through logistic regression, modeling the probability of an event occurring on a particular day, combined with an arbitrary fixed effect that measures the probability of any event occurring on that day. Direct estimation of fixed effects, however, generates bias in logistic regression (Greene 2004). Consistent estimation can instead be performed through conditional logistic regression, which estimates the model by conditioning on the number of events that occur within a fixed effect group. A conditional logistic model with a fixed effect for each discrete time unit is equivalent to the Cox proportional hazards model.⁵

⁵In additional, because the estimation procedure works by conditioning on the number of events

In addition, the use of the conditional logistic framework allows me to control for two important confounds for the estimation strategy described here: (1) the diversity of trader behavior, and (2) the diversity of currency characteristics.

Trader Diversity

Traders on the platform differ substantially in their trading characteristics, particularly in the amounts of money they can invest in and are willing to risk. On every distribution, from trading activity per day, to amount traded, to number of friends, some fraction of traders vastly exceeded their peers. At the same time, trader behavior greatly varied from day-to-day. Different traders followed different intra-week trading patterns, reacted to wins and losses by ramping their activity up or down, even exiting the market for months at a time or indefinitely.

This diversity is problematic to the extent that observed adoption or influence behavior is confounded with the characteristics of traders making the adoptions. That is, any observed result may be due to the fact that gregarious traders are also more eager to adopt currency pairs, generating a spurious peer effect. As such, one class of models I run gives each trader their own baseline hazard of adoption – that is, I condition on a fixed effect for each trader/day. Estimation then relies on the distribution of peer influence across the currency pairs a trader is at risk of adopting on a given day.

Currency Pair Diversity

Just as traders vary substantially in their characteristics, so do currency pairs. Different currency pairs have different trading characteristics, ranging from typical volatility

that occur within a group, the model ignores groups in which no events occurred, allowing them to be dropped from estimation. This allows me to reduce the effective number of observations needed to estimate any given model, greatly reducing the computational complexity of the problem.

to typical spreads and volume of trade, which derive from the characteristics of the currencies' home economies and governments, as well as the behavior of related currency pairs. At the same time, the trading characteristics and trading activity in particular currency pairs varies significantly from day-to-day, with much activity occurring around policy or news announcements. As such, the desirability of use for any two currency pairs can differ greatly each day.

Because of this, I expect that part of the observed adoption effect may be driven by external trends in currency desirability, so that we may observe two peers trading the same pair not because of influence between each other but because of external shocks. I attempt to control for this variability in two ways. As such, I estimate a second class of models by conditioning on a fixed effect for each currency-pair/day. The estimation here then relies on the distribution of peer influence across the various traders at risk of adoption on any given day.

Model

As described, I estimate two classes of model, one with fixed effects for trader/day, and one with fixed effects for currency-pair/day. Unfortunately, it does not appear possible to estimate both effects simultaneously, so I rely on consistency across the models to establish the results. The fixed effect procedure allows for a sparse model that avoids the need for extensive controls. The only variables included are those of immediate interest. Each observation in the data has a focal currency pair, the particular currency pair that the trader is at risk of adopting. Each trader can be at risk of adopting 133 different pairs on any given day – each is represented by a distinct observation. Because traders differed in their levels of activity and number of friends, distributions of peer activity and outcomes varied substantially. To reduce

sensitivity to outliers, the model relies on dummy variables to examine the effects of peer influence instead of testing for more specific functional forms across the observed range of values.

Variables of Interest

Alter Trade The question of interest is whether experiencing social influence produces consideration of a focal currency. This is a binary variable encoding whether *any* of a trader’s friends traded in the focal currency in the prior 14 days.

Alter Profit To control for the possibility that traders are influenced by the profitability of their friends’ trades, I construct an additional variable measuring whether the majority of peer trades in a focal currency pair made money. This is a binary variable encoding whether alters’ trades in the focal currency in the previous 14 days were more profitable than not. Alternative constructions of this variable that compared the ratio of profitable to losing trades against the community average or alter average showed similar patterns of results.

Currency Rank Estimation of this model relies on the estimation of interaction effects between measures of peer influence and the characteristics of the focal currency pair. There is no choice but to assume some parametric form for this interaction, even for models conditioning on unique baseline rates of adoption for each pair. This variable encodes the currency pair’s popularity rank on the platform: top currency EUR/USD has rank 1; the XAU/CNY pair, with only one observed trade, has rank 133. In accordance with hypothesis 3, we would expect currency pairs of higher rank to experience a greater effect of social influence. For robustness, I construct and report

several variants of this variable: I report results for the logarithm of currency pair rank as well as a dummy variable indicating whether the rank was greater than 10. Alternate functional forms substantiate the reported results.

Additional Controls

Time varying characteristics of both traders and currency pairs complicate the analysis of adoption. While trader or pair fixed effects can eliminate one source of confounding, the other source must be controlled for parametrically.

One potential problem is that veteran, popular, traders may both experience more social influence and be less susceptible to it. This biases estimates of the peer influence effect downward. While trader fixed effects can remove this source of confounding, currency pair fixed effects do not. As such, in some models I control for the size of a trader’s peer group.

Friends This measures the number of friends a trader has made up to a given trading day. Because the distribution of friend group size is highly skewed, I take the logarithm of the number of friends.

Results

Table 1 presents the results of conditional logistic regression on the probability of adoption, conditioning out a unique baseline rate of adoption for each trader. The coefficients here can be interpreted as the logarithm of a multiplier of the unobserved baseline rate, analogous to a Cox proportional hazard model. Negative coefficients

	(1) Baseline	(2) Log Rank	(3) Rank > 10
Alter Trade	2.508*** (0.023)	-1.446*** (0.043)	0.156*** (0.027)
Alter Profit	0.098*** (0.025)	0.011 (0.050)	0.080** (0.030)
Rank		-1.157*** (0.004)	-2.826*** (0.010)
Trade*Rank		0.971*** (0.021)	2.383*** (0.046)
Profit*Rank		0.024 (0.025)	-0.065 (0.055)
AIC	437216.317	340489.994	373876.040
Num. events	48961	48961	48961
Num. obs.	5001997	5001997	5001997

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 1: Conditional Logit (Trader/Day): Chance of Currency Adoption

represent a lower-than-baseline chance to adopt a currency pair; positive coefficients a greater chance.

Model 1 tests for an overall effect of peer influence in the setting. As predicted, controlling for trader fixed effects, we see a strong positive influence of peer use. The coefficient for *Alter Trade* implies that observing an alter trading in a particular currency pair makes a trader 12 times ($\exp(2.508) = 12.28$, 1100 percent) more likely to adopt a currency pair on a given day. The coefficient on *Alter Profit*, though significant, is much smaller in magnitude than the coefficient for trading, implying that observing profits only increases the chance of adoption by a factor of 1.1 ($\exp(0.098) = 1.10$, 10 percent) above and beyond the pure influence effect. These results support Hypotheses 1 and 2.

Model 2 adds an interaction with the logarithm of currency pair rank. As expected the

main effect of currency rank is highly negative – rarer currency pairs are less likely to be adopted, with the rarest pair being only 0.3 percent as likely to be adopted as the most common. Though the main effect of *Alter Trade* becomes significantly negative, its interaction with rank is positive, such that the overall effect of influence becomes positive at pairs of rank greater than 4. Moreover, there is neither a main effect nor an interaction of trade profitability. These results once again uphold hypotheses 1 and 2, and in addition, support hypothesis 3.

Model 3 presents an alternative specification of rank: here rank is a dummy variable indicating whether a currency’s rank is greater than 10. These results confirm those of model 2. In addition, the main effect of *Alter Trade* is slightly positive, indicating that even the top 10 currency pairs are subject to influence (17 percent), and that the negative main effect in model 2 is an artifact of the functional form.

As discussed above, the trader fixed effects of table 1 do not control for the time-varying desirability of currency pairs. Table 2 presents estimates of conditional logistic regression conditioning out a baseline hazard for each currency pair instead. Models 4 through 6 present analogous estimates to models 1 through 3. The positive interaction effect between *Alter Trade* and *Rank* once again supports hypothesis 3. The main effect of *Alter Trade* is negative, however, and the overall influence effect in model 5 only becomes positive for pairs of rank greater than 14. As discussed above, though, this estimate may be biased downwards if popular traders are also less likely to be influenced.

Model 7 adds an interaction for the size of a trader’s peer group to model 4. This model indicates that the negative peer influence of model 4 is biased by the un-susceptibility to influence of popular traders. The model suggests that traders are much less likely to adopt currency pairs in the presence of larger peer groups, but the lack of an

interaction on $Trade*Friends$ suggests that even popular traders are no less influenced by their peers. Model 8 elaborates this story by including an interaction with the logarithm of currency pair rank. Here, we once again see a strong positive interaction on $Trade*Rank$, indicating that all traders are more influenced by their peers for rare pairs than for common pairs. Large peer groups make traders less likely to adopt currency pairs overall, but do not strongly affect the magnitude of peer influence – all interactions are non-significant and small. Figure 3 plots the predicted effect of model 8, highlighting the effect of influence on traders with peer group sizes at the median (2) and the 95th percentile (42). The model indicates that the effect of peer influence becomes positive for pairs with rank greater than 3.

	(4) Baseline	(5) Log Rank	(6) Rank > 10	(7) Baseline	(8) Log Rank	(9) Rank > 10
Alter Trade	-0.236*** (0.020)	-0.661*** (0.046)	-0.350*** (0.024)	0.027 (0.055)	-0.376** (0.120)	-0.073 (0.065)
Alter Profit	0.028 (0.024)	0.056 (0.054)	0.065* (0.028)	-0.005 (0.066)	0.026 (0.144)	0.025 (0.078)
Trade*Rank		0.246*** (0.024)	0.427*** (0.045)		0.290*** (0.062)	0.632*** (0.123)
Profit*Rank		-0.009 (0.028)	-0.133* (0.054)		-0.044 (0.075)	-0.307* (0.152)
Friends				-0.110*** (0.004)	-0.186*** (0.012)	-0.147*** (0.007)
Trade*Friends				-0.019 (0.019)	0.012 (0.045)	-0.001 (0.024)
Profit*Friends				0.017 (0.022)	0.033 (0.052)	0.023 (0.027)
Rank*Friends					0.034*** (0.005)	0.070*** (0.009)
Trade*Rank*Friends					-0.031 (0.022)	-0.105** (0.040)
Profit*Rank*Friends					0.003 (0.026)	0.047 (0.048)
AIC	702292.406	702036.846	702151.156	701492.967	701175.482	701278.627
Num. events	48961	48961	48961	48961	48961	48961
Num. obs.	59495273	59495273	59495273	59495273	59495273	59495273

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 2: Conditional Logit (Currency Pair/Day): Chance of Currency Adoption

Figure 3: Effect of social influence moderated by trader's number of friends

Finally, model 9 presents an alternative specification of model 8, with a dummy variable indicating currency pairs with rank greater than 10. These results largely confirm those of model 8. The main effect of *Alter Trade* becomes non-significant and close to 0, and the negative interaction of *Profit*Rank* suggests that traders may be less likely to adopt rare profitable currency pairs than rare winners. This result is puzzling, but imprecisely estimated. The positive *Trade*Rank* interaction once again confirms that social influence becomes stronger for rare currency pairs. Overall, this set of models confirms hypothesis 3, while providing weaker evidence for hypotheses 1 and 2.

The fact that negative peer influence effects attenuate as models more precisely control for trader characteristics suggests that there may be additional sources of trader-level

variation that could be suppressing the peer influence effect. Without a strong model of trader behavior, any attempt to construct proxy variables for these characteristics would be speculative. As a robustness check, additional models that conditioned out trader fixed effects and controlled for time-varying effects of currency pair desirability confirmed the results of models 1 through 3, once again suggesting that the data supports the three hypotheses.

Discussion

The findings confirm all three of the study’s hypotheses. First, traders were more likely to adopt currency pairs after observing peer trading in those pairs. Second, this jump was caused by the observation of trading behavior, not by the observation of profitable trades. Finally, this social influence effect is stronger for rare currency pairs than it is for common pairs, indicating that social influence is filling in gaps in pre-existing attention to currency pairs, which differ on account of external exposure.

This process takes place in a setting in which existing theories of diffusion should predict no effects: traders do not discover new trading options from peers, and traders do not discover the value of existing options. The setting serves to isolate consideration as the mechanism behind the observed pattern of diffusion. It does not, however, serve to explain precisely what happens as consideration changes. In particular, there are several competing accounts for how a consideration process might work, and they deserve discussion.

Two Sources of Consideration

This paper argues that social influence causes traders to pay attention to and consider objects they otherwise would not have. This occurs because influence changes the perceived attributes of the object that render it relevant to a decision. There are two major competing accounts for how this process of consideration plays out in a trader's head: First, the changed attributes may belong to schemas of appropriate trading behavior – certain currency pairs that were previously considered unworthy now become appropriate trading instruments. Second, the changed attributes may proxy for cognitive state – influence exposes traders to particular objects, making them more cognitively salient and easier to access in future decisions (Tversky and Kahneman 1973; Bargh, Chen, and Burrows 1996). Since both mechanisms describe internal cognition processes, they are difficult to tease apart in observational data on behavior. Still, theory on cognitive schemas suggests a weak test to distinguish the two mechanisms in this setting.

Trading behavior may be driven by schemas of appropriate behavior in the domain. Such schemas may define trading patterns and situations, the interactions between particular currencies in those situations, and the correct currencies and pairs to use to exploit those situations. These schemas need not be correct. Research on interpersonal influence and persuasion, as well as on cognitive schemas, suggests that levels of prior expertise in a domain moderate how receptive individuals are to persuasion attempts and new information. In particular, though expert schemas of domains are larger than those of novices, they are more stable and less susceptible to change in light of new information (Dane 2010). On the other hand, though experts may be more capable of responding to information, they may have less motivation to do so if the information is deemed unimportant, leaving them vulnerable to non-conscious perceptual cues

in influence situations (Petty and Cacioppo 1986). As such, if influence operates by changing trading schemas and not through cognitive salience alone, we expect experts to be less likely to change their behavior because of influence than novices on the site. On the other hand, if influence primarily affects cognitive salience, we would expect no difference between experts and novices.

As it happens, the platform has an expert/novice distinction within its user base. Part of the platform’s strategy is to allow traders to follow the trades of company-selected ‘leaders,’ automatically copying the trades those leaders make. These leaders tend to be more established and more experienced traders in the foreign exchange community, whereas other traders are more likely to be newcomers. If consideration operates through a schematic change mechanism, we should expect new traders to be particularly receptive to social influence and leaders to be particularly un-receptive. As such, we can make two further predictions, extensions of hypotheses 1-3:

Hypothesis 4 Leaders will be less likely to adopt new currency pairs due to social influence than non-leaders.

Hypothesis 5 Leaders will not be more influenced by rarer currency pairs than by common currency pairs.

I test these hypotheses as interaction effects of leader status with the models described in table 1 – table 3 presents results of this analysis. Model 10 once again finds a positive *Alter Trade* effect and a much smaller *Alter Profit* effect, and it finds a small negative interaction of the influence term with leader status. This seems to support

hypothesis 4. To test hypothesis 5, model 11 includes interactions with the logarithm of currency pair rank. Here, all interactions with trade leader status disappear. The magnitude of the *Alter Trade*Leader* interaction remains roughly similar in magnitude to the previous model, though much less precise. Figure 4 presents the predicted effects of model 11, comparing the effect of observing peer influence on leaders and regular traders, as compared to a regular trader experiencing no influence in a focal currency pair. The figure includes non-significant effects. Finally, model 12 replaces the logarithm of rank by a binary variable indicating whether a pair's rank is greater than 10. This model once again suggests a weak negative interaction between leader status and the peer influence effect. Overall, hypothesis 4 is at best weakly supported and hypothesis 5 not at all.

	(10) Baseline	(11) Log Rank	(12) Rank > 10
Alter Trade	2.519*** (0.023)	-1.438*** (0.043)	0.166*** (0.027)
Alter Profit	0.089*** (0.026)	0.012 (0.050)	0.076* (0.030)
Alter Trade*Leader	-0.389** (0.143)	-0.534 (0.321)	-0.406* (0.185)
Alter Profit*Leader	0.289 (0.148)	0.201 (0.355)	0.175 (0.196)
Rank		-1.158*** (0.004)	-2.830*** (0.010)
Trade*Rank		0.973*** (0.021)	2.386*** (0.047)
Profit*Rank		0.018 (0.025)	-0.086 (0.056)
Rank*Leader		0.049* (0.021)	0.150* (0.062)
Trade*Rank*Leader		0.066 (0.133)	0.057 (0.258)
Profit*Rank*Leader		0.054 (0.151)	0.354 (0.297)
AIC	437212.539	340478.239	373857.664
Num. events	48961	48961	48961
Num. obs.	5001997	5001997	5001997

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 3: Conditional Logit (Trader/Day): Effects of Leader Experience

Figure 4: Effect of social influence on leaders and regular traders

The evidence here suggests that leaders are no less susceptible to social influence than are regular traders. To the extent that this setting follows the prior literature indicating that expertise stabilizes schemas, this result is weak evidence against the argument that social influence primarily affects traders' schemas of appropriate behavior in the setting. Instead, it appears to indicate that influence works through an implicit cognitive decision process. Both leaders and regular traders are equally influenced by the actions of their peers. It remains for further research to illuminate whether this is a failure of the data to be informative, or whether some feature of the foreign exchange market makes expertise less durable in this setting.

General Discussion

There are two further points that merit discussion and future work. First, one of the drawbacks of the setting is the inability to reliably detect discontinuance behavior. One of the implications of a consideration-driven adoption model is that continuing use of an object depends on continuing consideration of it. While it seems strange to believe that objects can leave an information set or change value in the absence of some structural change, attention may be a more fleeting phenomenon. Attention serves to constrain the set of objects that require evaluation, and limited attention can be beneficial when evaluation is costly (March and Simon 1958). The persistence of attention depends on the changing situation of the agent and may require constant reinforcement. To the extent that consideration is driven by a non-conscious exposure process, we may expect to observe discontinuance behavior in the absence of reinforcing social influence. Similarly, we may observe discontinuance if people observe the discontinuance of their peers and take this as evidence of changing standards. While it may be possible to look for evidence of such an effect in the foreign exchange market, lack of use there may indicate the absence of favorable situations for currency pair use. A more natural setting would feature objects whose continual use and disuse is less ambiguous. Such attention-driven discontinuance could account for phenomena of discontinuance behavior such as declining rates of condom use in HIV-susceptible populations (Disease Control and Prevention 2012).

Second, identification of peer influence in this study suffers from the well-documented problem of homophily (Aral, Muchnik, and Sundararajan 2009; Shalizi and Thomas 2010). In this context, it is unclear to what extent traders are simply making friendships with those who are predisposed to trading in the same currency pairs. In practice, this does not appear to be a serious concern – though traders use currencies in clusters,

they do not often appear to adopt them in clusters. The gap between adoption times of a focal trader and his peers can be on the order of months. In addition, homophily on trader susceptibility to learn about a currency is not in itself theoretically problematic here – traders may well be predisposed to being interested in a particular currency pair but only adopt it conditional on being able to learn about it from peers. Nevertheless, this study would benefit from the ability to identify the effect of homophily by finding some way to instrument for tie formation. No such instrument presents itself.

Conclusions

This study demonstrates an elaboration of the typical mechanisms of diffusion, establishing an intermediate step between the diffusion of information about an object and the influence-driven change in valuation that favors the object’s adoption. Social influence acts to direct attention to an object, allowing it to enter the consideration set for a decision. Existing models of diffusion typically assume the underlying value of a given object and assume that where its diffusion isn’t slowed by its poor quality, lack of information accounts for failure of adoption. Instead, this study finds that diffusion of object consideration can account for the success or failure of a diffusion process.

More broadly, this study integrates existing theory of diffusion processes into a single model of the object adoption decision. Existing research on informational and valuational effects of social influence emerge as stages in a sequential decision process, and the model reveals additional stages of the process that are novel to the literature. Without knowledge of an object, there is no consideration of it. Without consideration, there is no evaluation. The stages of the decision process govern which

patterns of influence can affect the adopter. If an object is out of consideration, diffusion processes and networks structures that affect its value will have no effect on the adoption outcome. Use of the unified adoption model can help diagnose the course of diffusion processes in the wild, determining which stages of the decision process served to accelerate or deter adoption. Furthermore, the model can help predict what features of a new object are most critical to its adoption, whether adoption relies most on the diffusion of information, attention, or value.

This study also answers and echoes the call by Strang and Soule (1998) for more comparative analysis of diffusion. Variance in object characteristics within a setting can reveal why certain aspects of social structure were responsible for diffusion. Conversely, understanding the social structures involved in a particular setting can reveal the content of the influence process, shedding light on whether structures more conducive to information flow, consideration, or value-creation were at work.

Finally, this study opens an opportunity for the integration of research on diffusion and the advancing work on the categorical determinants of decision-making. On the one hand, diffusion of categories (cf. Rossman 2014) may determine the viability of adoption for a category member by opening space for object consideration. On the other hand, the spread of categories themselves can be subject to diffusion modeling. To date, work on the diffusion of conceptual categories has been much more limited than work on the spread of objects, and it is unclear how similar theories of the two processes may be.

Far from simply providing channels for the flow of knowledge, or environments for the production of network externalities, social structures shape local attention and perception of diffusing objects. Key actors in the social structure may be responsible not just for disseminating information but shaping their peers' understanding of a

product or practice – how it is to be used, when, by whom. Understanding how attention diffuses will go a long way to accounting not only for why some products manage to take off and others fail, but it will shed light on how objects can gain new meanings and transform in their journey through a population.

Appendix: Evidence Against Trader Skill

One of the key identifying assumptions of this paper is that traders do not adopt the behavior of peers because of peers' underlying skill in identifying good trading opportunities. That is to say, peers provide focal traders not with evaluative information about the inherent profitability of a particular currency pair at a particular time, but with normative information about which currency pairs are worth consideration. As described above, the major source of evidence behind this assumption comes from existing literature on trading and currency markets suggesting that traders almost always to be unable to beat the market and make reliable returns.

This appendix presents internal evidence for the same claim. Emulating Barber et al. (2014), I test whether traders' past returns predict their future returns. Trader skill should generate serial correlation of trading returns in the population. If there is evidence that trader skill exists, then members of the site may be able to detect such skill in a way undetectable to the researcher, so that their adoption behavior is (potentially) rationally driven by attempts to emulate the returns of high performers. If not, then adoption behavior may be driven either by superstitious learning or by the attention effects proposed in this paper.

I measure trader performance every month for the five most popular currency pairs on the platform: EUR/USD, GBP/USD, AUD/USD, USD/JPY, and EUR/JPY. Table [A1](#) presents OLS estimates of the regression of current month performance on prior month performance by currency pair. As the table suggests, there appears to evidence of skill in four of the five currencies. Digging deeper, however, I partition traders into quantiles of skill: each month, I sort all traders by their returns in the previous month and regress their current month performance on their quantile. Table [A2](#) shows

the results of this test: ‘%ile: 0-1’ represents the worst traders, while ‘%ile: 99-100’ represents the top traders. Here, it becomes clear that evidence of trader skill appears to be driven almost entirely by the downside – while the best traders don’t appear to be able to replicate their results month to month, the worst traders are reliably bad.

	EURUSD	GBPUSD	AUDUSD	USDJPY	EURJPY
Prior Month Return	0.0000 (0.0052)	0.0313*** (0.0075)	0.0903*** (0.0079)	0.1541*** (0.0093)	0.0821*** (0.0111)
(Intercept)	−0.1210 (0.0697)	−0.0251*** (0.0029)	−0.0369*** (0.0057)	−0.0342*** (0.0063)	−0.0441*** (0.0075)
Adj. R ²	0.0000	0.0007	0.0070	0.0171	0.0044
Num. obs.	37055	22985	18458	15749	12002

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A1: Predictability of Trader Profits: Month-to-Month Returns

These results are not unexpected. Interviews with currency traders suggest that one of the major skills traders have to learn is discipline to a program: inexperienced traders tend to take winning positions early and leave losing positions open long. Evidence from the platform confirms this notion: traders on average experience more losses than gains, but their losses are bigger, such that on average traders lose money. As such, trader psychology appears to make it much easier to be a low-skilled trader than to break even.

This evidence confirms findings of other researchers about the unviability of profitable personal trading. Skilled traders appear not to exist on this platform, suggesting that the assumption that traders do not attempt to emulate the latent success of their peers holds in this context, and that we can reliably interpret adoption behavior as incidences of superstitious learning or learning over what objects belong in the consideration set for trading.

	EURUSD	GBPUSD	AUDUSD	USDJPY	EURJPY
%ile: 0-1	−0.3199 (0.6195)	−0.1283*** (0.0285)	−0.2733*** (0.0595)	−0.7644*** (0.0597)	−0.1828 (5.9191)
%ile: 1-5	−0.1448 (0.2963)	−0.0573*** (0.0134)	−0.0504 (0.0280)	−0.0708* (0.0277)	−0.0821 (2.7192)
%ile: 5-20	−0.0639 (0.1528)	−0.0345*** (0.0069)	−0.0532*** (0.0144)	−0.0417** (0.0142)	−0.0352 (1.3945)
%ile: 20-40	−0.0395 (0.1324)	−0.0293*** (0.0060)	−0.0549*** (0.0125)	−0.0449*** (0.0123)	−0.0410 (1.2073)
%ile: 40-60	−0.3315* (0.1323)	−0.0211*** (0.0060)	−0.0309* (0.0125)	−0.0310* (0.0123)	−0.0479 (1.2075)
%ile: 60-80	−0.0537 (0.1325)	−0.0266*** (0.0060)	−0.0371** (0.0125)	−0.0146 (0.0123)	−0.0549 (1.2080)
%ile: 80-95	−0.0219 (0.1534)	−0.0212** (0.0069)	−0.0016 (0.0144)	0.0101 (0.0142)	3.5884* (1.3990)
%ile: 95-99	0.0064 (0.2965)	−0.0167 (0.0134)	0.0044 (0.0280)	0.0074 (0.0278)	0.0269 (2.7325)
%ile: 99-100	0.1862 (0.5453)	−0.0026 (0.0237)	0.0232 (0.0494)	−0.0394 (0.0465)	0.0428 (4.4892)
Adj. R ²	0.0000	0.0043	0.0028	0.0094	−0.0002
Num. obs.	43682	28343	22785	20219	15621

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A2: Predictability of Trader Profits: Quantile of Prior Returns

References

- Abbey, Boris S., and John A. Doukas. 2014. "Do individual currency traders make money?" *Journal of International Money and Finance*, October. Elsevier Ltd, 1–20. doi:[10.1016/j.jimonfin.2014.10.003](https://doi.org/10.1016/j.jimonfin.2014.10.003).
- Antonio, José Rosa, Joseph F. Porac, Jelena Runser-Spanjol, and Michael S. Saxon. 1999. "Sociocognitive Dynamics in a Product Market." *Journal of Marketing* 63: 64–77.
- Aral, Sinan, and Marshall Van Alstyne. 2011. "The Diversity-Bandwidth Trade-off." *American Journal of Sociology* 117 (1): 90–171.
- Aral, Sinan, and Dylan Walker. 2012. "Identifying influential and susceptible members of social networks." *Science* 337 (6092): 337–41. doi:[10.1126/science.1215842](https://doi.org/10.1126/science.1215842).
- Aral, Sinan, Lev Muchnik, and Arun Sundararajan. 2009. "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks." *Proceedings of the National Academy of Sciences of the United States of America* 106 (51): 21544–9. doi:[10.1073/pnas.0908800106](https://doi.org/10.1073/pnas.0908800106).
- Banerjee, Abhijit, Arun G. Chandrasekha, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science*, NBER, 341 (6144): 1–56.
- Barber, Brad M., and Terrance Odean. 2000. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." *Journal of Finance* LV (2): 773–806.
- Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2014. "The cross-section of speculator skill: Evidence from day trading." *Journal of Financial Markets* 18 (March). Elsevier: 1–24. doi:[10.1016/j.finmar.2013.05.006](https://doi.org/10.1016/j.finmar.2013.05.006).
- Bargh, John A., Mark Chen, and Lara Burrows. 1996. "Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Activation on Action." *Journal of Personality and Social Psychology* 71 (2): 230–44.
- Bessembinder, Hendrik, and Kalok Chan. 1998. "Market Efficiency and the Returns to Technical Analysis." *Financial Management* 27 (2): 5–17.
- Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113 (3): 702–34.
- Dane, Erik. 2010. "Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective." *Academy of Management Review* 35 (4): 579–603. doi:[10.5465/AMR.2010.53502832](https://doi.org/10.5465/AMR.2010.53502832).
- Davis, Gerald, and Henrich Greve. 1997. "Corporate Elite Networks and Governance Changes in the 1980s." *American Journal of Sociology* 103: 1–37.

- Disease Control, Centers for, and Prevention. 2012. *Trends in HIV-Related Risk Behaviors among High School Students—United States, 1991-2011*. 29. *Morbidity and Mortality Weekly Report*. Vol. 61. Centers for Disease Control; Prevention.
- Edelman, Lauren B., Christopher Uggen, and Howard S. Erlanger. 1999. “The Endogeneity of Legal Regulation: Grievance Procedures as Rational Myth.” *American Journal of Sociology* 105 (2): 406–54.
- Granovetter, Mark S. 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78 (6). University of Chicago Press; JSTOR: 1360–80. doi:[10.1086/225469](https://doi.org/10.1086/225469).
- Greene, William. 2004. “The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects.” *Econometrics Journal* 7: 98–119.
- Hannan, Michael T., László Pólos, and Glenn R. Carroll. 2007. *Logics of Organization Theory: Audiences, Codes, and Ecologies*. Princeton: Princeton University Press.
- Heimer, Rawley Z, and David Simon. 2012. “Facebook Finance: How Social Interaction Propagates Active Investing.”
- International Settlements, Bank for. 2013. *Triennial Central Bank Survey: Foreign exchange turnover in April 2013: preliminary global results*. April. Bank for International Settlements.
- Jackson, Matthew O. 2010. *Social and Economic Networks*. Princeton: Princeton University Press.
- Kahneman, Daniel, and Shane Frederick. 2005. “A Model of Heuristic Judgment.” In *The Cambridge Handbook of Thinking and Reasoning*, 267–94. New York: Cambridge University Press.
- March, James G. 1994. *A Primer on Decision Making: How Decisions Happen*. New York: Free Press.
- March, James G., and Johan P. Olsen. 1976. “Organizational Choice under Ambiguity.” In *Ambiguity and Choice in Organizations*. Universitetsforlaget.
- March, James G., and Herbert A. Simon. 1958. “Cognitive Limits on Rationality.” In *Organizations*.
- Menkhoff, Lukas, and Mark P. Taylor. 2007. “The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis.” *Journal of Economic Literature* 45 (4): 936–72.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. 2012. “Currency momentum strategies.” *Journal of Financial Economics* 106 (3). Elsevier: 660–84. doi:[10.1016/j.jfineco.2012.06.009](https://doi.org/10.1016/j.jfineco.2012.06.009).
- Muchnik, Lev, Sinan Aral, and Sean J. Taylor. 2013. “Social Influence Bias: A Randomized Experiment.” *Science* 341 (August): 647–51.

- Pan, Wei, Yaniv Altshuler, and Alex (Sandy) Pentland. 2012. “Decoding Social Influence and the Wisdom of the Crowd in Financial Trading Network.” In *2012 International Conference on Social Computing*, 203–9. Ieee. doi:[10.1109/SocialCom-PASSAT.2012.133](https://doi.org/10.1109/SocialCom-PASSAT.2012.133).
- Park, Cheol-Ho, and Scott H. Irwin. 2007. “What Do We Know About the Profitability of Technical Analysis?” *Journal of Economic Surveys* 21 (4): 786–826. doi:[10.1111/j.1467-6419.2007.00519.x](https://doi.org/10.1111/j.1467-6419.2007.00519.x).
- Park, Cheol-ho, and Scott H. Irwin. 2010. “A Reality Check On Technical Trading Rule Profits in the U.S. Futures Markets.” *Journal of Futures Markets* 30 (7): 633–59. doi:[10.1002/fut](https://doi.org/10.1002/fut).
- Petty, Richard E, and John T Cacioppo. 1986. “The Elaboration Likelihood Model of Persuasion.” In *The Elaboration Likelihood Model of Persuasion*.
- Podolny, JM. 2001. “Networks as the Pipes and Prisms of the Market.” *American Journal of Sociology* 107 (1): 33–60.
- Roberts, John H., and James M. Lattin. 1991. “Development Testing of a Model of Consideration Set Composition.” *Journal of Marketing Research* 28 (4): 429–40.
- Rogers, Everett M. 2003. *Diffusion of Innovations*. 5th ed. New York: Free Press.
- Rossman, Gabriel. 2014. “The Diffusion of the Legitimate and the Diffusion of Legitimacy.” *Sociological Science* 1 (March): 49–69. doi:[10.15195/v1.a5](https://doi.org/10.15195/v1.a5).
- Schulmeister, Stephan. 2009. “Profitability of technical stock trading: Has it moved from daily to intraday data?” *Review of Financial Economics* 18 (4). Elsevier Inc.: 190–201. doi:[10.1016/j.rfe.2008.10.001](https://doi.org/10.1016/j.rfe.2008.10.001).
- Shalizi, Cosma Rohilla, and Andrew C. Thomas. 2010. “Homophily and Contagion Are Generically Confounded in Observational Social Network Studies.” *ArXiv*, April, 1–27.
- Simon, David. 2013. “Social Networks and Price Discovery.”
- Strang, David, and Sarah A. Soule. 1998. “Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills.” *Annual Review of Sociology* 24 (1998): 265–90.
- Tversky, Amos, and Daniel Kahneman. 1973. “Availability: A heuristic for judging frequency and probability.” *Cognitive Psychology* 5: 207–32. doi:[10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9).
- Zuckerman, Ezra W. 1999. “The Categorical Imperative: Securities Analysts and the Illegitimacy Discount.” *American Journal of Sociology* 104 (5).