

U-Shaped Conformity in Online Social Networks*

Monic Sun

Boston University

Email: monic@bu.edu

Xiaoquan (Michael) Zhang

Chinese University of Hong Kong

Email: zhang@cuhk.edu.hk

Feng Zhu

Harvard Business School

Email: fzhu@hbs.edu

August 1, 2018

*We are grateful to Binghao Cheng, Huijie Xing, Yongqing Xu, and Shuhua Zhou for their help with data collection. Financial support from the NET Institute is greatly appreciated. For helpful discussions and comments, we thank Carey Morewedge, Navdeep Sahni, and seminar participants at Arizona State, Boston U, London Business School, Penn State, U of Cambridge, U of Maryland, U of Minnesota, U of Toronto and the following conferences: Marketing Analytics and Big Data (Columbia), Big Data Marketing Analytics (Chicago), CODE (MIT), ISMS Marketing Science (Istanbul), NET Institute (Berkeley), Theory + Practice in Marketing (Evanston), UT Dallas FORMS, and the ZEW Conference on the Economics of Information and Communication Technologies.

U-Shaped Conformity in Online Social Networks

Abstract

We explore how people balance their needs to belong and to be different from their friends by studying their choices of a virtual-house wall color on a leading Chinese social-networking site. The setting enables us to randomize both the popular color and the adoption rate at the individual level so that our experimental design minimizes informational social influence, homophily, and group-identity signaling to the general public. We find that there exists significant social influence within a user's friend circles. While learning about the most popular color among a user's friends generally increases the likelihood for the user to adopt that color, conformity first decreases and then increases with the adoption rate of that choice, which ranges from 50% to 100%. In addition, users who are minority, newer, or of lower social-economic status are more likely to conform upon learning about the popular choice. Our findings are consistent with optimal distinctiveness and middle-status conformity theories and have implications for designing normative marketing campaigns.

Keywords: conformity, normative social influence, social networks, field experiment

1. Introduction

A number of studies in social psychology (e.g., Asch 1952, Goffman 1959, Aronson 1972), economics (e.g., Bernheim 1994, Akerlof and Kranton 2000) and marketing (e.g., Venkatesan 1966, Miniard and Cohen 1983) have documented conformity toward popular choices in one's offline social network. The rise of social networking websites today has provided their users with unprecedented opportunities to be connected.¹ The new context raises three questions related to conformity that we explore in this paper. First, does there exist significant social influence among a user's friends in the context of social networking sites? Second, how do the users' reactions to the majority choice among their friends change with the popularity of the choice? Third, how do people with different demographic and network characteristics react differently to the popular choice among their friends?

To answer these questions, we conduct a large-scale field experiment on a leading social networking site (SNS) in China, Kaixin001.com (henceforth referred to as Kaixin), using a built-in game on the site, Virtual Homes. Participants can paint their virtual houses freely with a color of their choice. Participants are randomly allocated into different experimental conditions in which they are exposed to different types of information regarding the recently popular colors among their friends. They are then offered a choice to either repaint the wall or continue with their previous choice. To minimize confounding factors to the identification of causal social influences, such as homophily, both the popular color and the associated adoption rate are randomly generated.² We also collect demographic information from their registration pages and network-

¹ Online social networking platforms surpassed 2 million users in 2016 and as of 2017 and daily social media usage of global Internet users amounted to 135 minutes per day. Source: <https://www.statista.com/statistics/433871/daily-social-media-usage-worldwide/>, accessed in July 2018.

² Homophily refers to the phenomenon that friends often have similar tastes to begin with, which may confound causal inferences on social influences (e.g., Currarini and Pin 2009; Kossinets and Watts 2009).

related information such as their numbers of friends. In our analysis, we focus on whether a user repaints his virtual house with the popular color among his friends and examine how such conforming behavior varies with the adoption rate of the choice and the characteristics of the participant himself.³

We obtain three sets of main findings. First, there exists social influence within the users' friend circles, in the sense that their overall likelihood of choosing a particular color often changes once that color is announced to be the most popular among their friends. This finding confirms our intuition that people react to their friends' choices. Second, as the adoption rate of the majority choice increases, a user's likelihood of conforming to that choice first decreases and then increases. In particular, a user is the least likely to conform when the adoption rate of the majority choice is around 80% and the most likely to conform when the adoption rate is near 50% or 100%. To our best knowledge, our study is the first to causally identify a U-shaped conformity curve with a large-scale social network from the real world. While we do not have direct process evidence because our study is based on observations of changes in behavior instead of self-reported attitudes, the result seems to be consistent with the optimal distinctiveness theory (Brewer 1991, 1993), which posits that individuals are most likely to identify with moderately inclusive rather than extremely inclusive or exclusive groups (Leonardelli, Pickett and Brewer 2010).

Finally, we find that certain user characteristics have a significant impact on conformity. In particular, users who are less experienced, older in age, male, born in less affluent towns, or have fewer friends are on average more likely to conform. Given that the majority of Kaixin users are young and female, our results suggest that users of minority, newer, and lower social-economic status are more likely to conform. This result is consistent with existing social psychology studies

³ We refer to the user as "he" in the paper for ease of exposition. There are both female and male users on the site.

that find that weakly identified members of a social network tend to conform to group norms to gain acceptance, while well-established members are secure enough that their main concern has shifted from gaining acceptance to signaling personal identity through demonstrating divergence (e.g., Goffman 1963, Aronson 1972).

Our results have interesting implications for marketing practitioners. First, our results suggest that even though people may behave differently online in general, there is still robust evidence of social influence. While one may think that friends connected through a social networking site may be physically far away from each other and the distance would reduce peer influence, we observe significant conformity toward popular choices among one's online friends. One potential reason is that on social networking sites, a user's actions can be quickly and widely observed by a large number of his friends and the benefits of conformity, such as improved communication with friends, may materialize faster.

Second, our findings suggest that a higher adoption rate is not always better in normative marketing campaigns. In recent years, normative campaigns have seen increasing popularity in the field of marketing, and they often prominently show the adoption rate of the desired behavior with the hope that people will conform to the desired behavior. For example, Dartmouth advertises that 74% of Dartmouth drinkers have fewer than four drinks on the average Friday night^{4,5} and

⁴ <http://www.dartmouth.edu/~healthed/focus/aod/norms.html>, accessed in March 2018.

⁵ According to the National Social Norms Center (<https://socialnorms.org/faqs/>, accessed in March 2018), normative campaigns have worked successfully in the domains of alcohol and tobacco use, middle-school bullying, seat belt use, driving under the influence, teen sexual behavior, parental behavior, seeking of mental health care, home energy use, littering, curbside recycling, voting, and tax compliance (e.g., Haines, Barker and Rice 2003, Gerber, Green and Larimer 2008, Perkins et al. 2010, Schultz 2010, Perkins, Craig and Perkins 2011). Other studies, at the same time, have documented the ineffectiveness of certain normative marketing campaigns around alcohol consumption, tobacco use, littering, and exercise behavior (e.g., Kallgren, Reno and Cialdini 2000, Cialdini 2003, Campo and Cameron 2006, Schultz et al. 2007, Hornik et al. 2008). Our study focuses on a less injunctive norm but the findings may still shed some light on why certain normative campaigns may be more effective than others.

Sensodyne advertises that 9 out of 10 dentists recommend their toothpaste.⁶ Our findings suggest that designers of such campaigns need to be mindful of the strength of the norm and the characteristics of the target audience.

1.1 Related Literature

Given the co-existence of the needs to belong and to be different, it is not surprising that many scholars have explored how people balance them (e.g., Brewer 1991, Hornsey and Jetten 2004).⁷ Brewer (1991, 1993) first proposes the optimal distinctiveness theory (ODT), which is an important milestone in this literature. It asserts that individuals desire to achieve an optimal balance of inclusion and distinctiveness within and between social groups and as a result, are most likely to identify with groups that are neither too inclusive nor too exclusive. In the context of our experiment, ODT would predict that a user would identify with a distinct subgroup of his friends.

Evidence for ODT in the sociology literature starts with Tajfel et al. (1971) and Leonardelli (2006), who find that people in general prefer to identify with groups of numerical minority. In addition, Lau (1989) finds that black respondents who lived in areas where 40-70% of the residents were black were significantly more likely to identify themselves as “particularly close” to blacks, Bearman and Brückner (2001) find that the likelihood for high school girls to commit to a public

⁶ <https://us.sensodyne.com/>, accessed in July 2018. The recommendation in this case may serve as a signal of high quality, and such signaling effect is minimized in our experiment to focus on the inter-personal social influence.

⁷ People often keep dissenting thoughts private for fear of social sanction (e.g., Deutsch and Gerard 1955), or take clearly wrong actions when such actions enable them to fit into the group (e.g., Asch 1952). They conform to the behavior of other in-group members (e.g., Chartrand, Rigotti and Rustichini 2007, Berger and Heath 2008) because of non-conscious mimicry (e.g., Chartrand and Bargh 1999), social pressure (e.g., Cialdini, Reno and Kallgren 1990, Austen-Smith and Fryer 2005, Mas and Moretti 2009) or the need to be consistent with their social (e.g., Benjamin, Choi and Strickland 2010) or group identities (e.g., Goette, Huffman and Meier 2006, Chen and Li 2009). On the other hand, Neo-Jamesian self theories propose that while people have various reasons to conform to others, they also have an innate need to feel different from others and that feeling some degree of uniqueness is generally essential to one’s well-being (Snyder and Fromkin 1980, Brewer 1991). Diverging from others is a meaningful way to establish one’s individual identity (e.g., Vignoles, Chrysoschoou and Breakwell 2000) and that people are more motivated to diverge from others when such divergence becomes more visible (e.g., Codol 1981). Importantly, being a nonconformist can have a positive valence in individualistic societies (e.g., Hofstede 1980, Triandis 1995) while being seen as immature in collectivistic cultures (e.g., Markus and Kitayama 1994, Kim and Markus 1999).

virginity pledge increased as the percentage of respondents increased until 40% of respondents had committed to such pledges, Abrams (2009) demonstrates that young adults with moderately exclusive music preferences reported the greatest degree of behavioral identification, and Leonardelli and Loyd (2016) find that the greater identification comes from a higher level of trust that stems from optimal distinctiveness of the group.⁸ Different from these studies, our study does not use self-reported attitudes and focuses instead on spontaneous choices of a large pool of users on a real social network. Consequently, the choice can be posited to be popular among a user's friends with mutual agreement rather than the general public and the adoption rate is randomized on a continuous scale to minimize the confounding effects of homophily.

Prior studies also suggest that minority members of a society may be more likely than others to conform to such norms. Lieberman and Mikelson (1995) document how African American parents are heavily influenced by mainstream culture at the time of naming their babies and Zhang, Zuckerman and Obukhova (2016) find evidence that Jewish immigrants deliberately choose established baby names so as to express their membership in U.S. society. Formally, Phillips and Zuckerman (2001) propose the middle-status conformity theory: high-status members feel confident and secure about their role incumbency while low-status members feel excluded regardless of their choices. Middle-status members, on the other hand, feel the most intense anxiety to belong and fear of disenfranchisement, and therefore are most likely to conform. They specifically point out that unless the lowest ranked candidates are observable outsiders, a simple negative relationship between status and conformity should be observed. Consistent with their prediction and our own findings, Jetten, Hornsey and Adarves-Yorno (2006) demonstrate that people are more motivated to portray themselves as conformist when they feel their intragroup

⁸ Bedea et al. (2010) find that the most distinctive group is preferred when the group's identity is threatened.

status to be relatively low, although they do not examine how conformity varies with popularity of the majority choice.

The remainder of the paper is organized as follows. We present the design of the field experiment in Section 2, data and analysis in Section 3, robustness checks in Section 4, and concluding remarks in Section 5.

2. Experimental Design

Founded in March 2008, Kaixin had more than 161 million registered users as of June 2014, roughly 10 percent of the Chinese population.⁹ According to a 2009 national study by the Chinese Internet Network Information Center, Kaixin is one of the most popular SNSs in China, hosting 26.4 percent of all SNS users there.¹⁰ Compared to other Chinese SNSs, Kaixin has superior coverage of urban users and appeals to white-collar office workers who love surfing the site at work, while its large user base manages to cover a variety of demographics. A typical user of Kaixin is 25 to 34 years old and has a college degree. Kaixin users are highly active: in August 2011, the site averaged more than 8 billion page views per week.¹¹ Most Kaixin users are registered with their real names, and two users must agree to a friendship before they can access each other's information as friends.

We conduct the experiment in the context of a built-in game on Kaixin, "Virtual Homes." Given its built-in nature, all registered users have automatic access to it, and bias due to game selection is therefore minimized. It is a non-competitive game in which a user can creatively customize the looks of his virtual house. He can, in particular, choose the color of the walls, from

⁹ Source: <http://www.kaixin001.com>, accessed March 2018.

¹⁰ Source: <http://www.web2asia.com/2009/11/12/latest-statistics-on-online-sns-usage-in-china>, accessed in March 2018. Other popular Chinese social networking sites are 51.com, QQ XiaoYou, Sina Space, and RenRen.

¹¹ Source: <http://tech.sina.com.cn/i/2011-08-11/14185915333.shtml>, accessed in March 2018.

among the following color choices: yellow, green, pink, blue, red, and gray. We provide an illustration of the game, translated from Chinese, in Figure 1.

We choose a random sample of active Kaixin users who used this application between January and March 2011 and had more than 10 friends on Kaixin as the subjects for our experiment. We require the minimum number of friends to increase the credibility of our experimental message, as it might be easy for a user with only a few friends to remember their friends' recent color choices and question the validity of our message. By the time when the user is presented the experimental message, he can no longer check his friends' color choices. To avoid confusing the users with inconsistent experimental messages, we make sure that each user can participate in our experiment only once: once a user has been selected into our experiment, he will not be selected again. Our experiment runs in two phases for a total of three weeks. After the first phase in January 2011, we check the data to ensure that all procedures were implemented correctly and then run a second phase in March 2011.

Prior to the experiment, when a user picks a color and confirms this choice by clicking the "Save" button, Kaixin records the color choice, applies it to the subject's virtual house, and displays "Successfully Saved!" in a pop-up window. The user then clicks an "OK" button to close the window. During the experiment, we first randomly assign each subject to one of two equally probable experimental conditions, A and B. The conditions differ by the experimental message that we insert into the pop-up window mentioned above. In condition A, the control group condition, no additional information is inserted. In condition B, the treatment condition, a user is shown the *randomly generated* most popular color supposedly chosen by his friends and that color's *randomly generated* adoption rate. To focus on conformity, the displayed popular color is always different from the subject's original choice. We also insert the following sentence to remind

the user that his choice will be visible to his friends: “Don’t forget to show your newly painted house to your friends.” Figure 2 shows an exemplary pop-up window that a subject can see in condition B (translated from Chinese), and Figure 3 displays a flowchart of the experimental design.¹²

For all the subjects, the customized pop-up window offers two buttons, “Repaint” and “OK,” with the first button leading the user back to the page of color choices and the second button confirming the color choice. We record all subjects’ choices of the wall color before and after they see the experimental messages, as shown by the timeline of events in Figure 4.

We take several measures to ensure that participants cannot easily detect that our messages are randomly generated. First, as mentioned above, we only include users with more than 10 friends. Second, to make the alleged adoption rate of the most popular color among one’s friends look plausible, we use the following procedure to generate the adoption rate: For a subject with n friends ($n > 10$), let x be the smallest integer such that $x/n \geq 1/2$.¹³ Generate a random integer m between x and n , and then use $[m/n]$ as the adoption rate. Given our procedure, for example, a subject with 20 friends would only be given adoption rates of 50 percent, 55 percent, 60 percent, and so on. The procedure makes it hard for subjects to question the adoption rate solely based on the number of friends they have on the site. Third, in all the experimental messages, we use a vague word, “Recently,” when giving information about the most popular color (see Figure 2). The subjects, being unsure over what time period we are claiming the color has been popular, are less likely to question the claim.

¹² We also allocate a group of users to Condition C who receive information on the most popular color on the entire Kaixin platform (GlobalColor). Figures 3 and 4 contain this condition and we discuss it in more detail in Section 4.

¹³ In Section 4, we discuss non-majority choices for which the adoption rate can be lower than 50%.

Our experimental setting allows us to overcome the common empirical challenges of identifying social influence: the difficulty of observing actual friendship structures and the existence of contextual and confounding factors (e.g., Manski 1993, Hartmann et al. 2008). Our use of familiar colors, for example, minimizes observational learning (e.g., Scharfstein and Stein 1990, Banerjee 1992, Bikhchandani, Hirshleifer and Welch 1992, Tucker and Zhang 2011) as all six colors are standard and the subjects can experiment with them before confirming their choices. Unlike studies of information cascades or herding, our study involves little to no uncertainty in the choices themselves and hence eliminates the subjects' need to infer the quality of these choices through observational learning. By randomly generating the most popular colors, we also minimize the confounding effects of homophily: the correlations between a subject's color preferences and those of his friends. Such correlations, if not carefully controlled for, would make it hard for us to make causal inferences regarding how one's choice is causally influenced, as opposed to merely being reflected, by his friends' choices.

In addition, the fact that a user's choice is observed only by friends on the site rules out his need to signal his group identity to non-friend users on the platform (e.g., Griskevicius et al. 2006, Berger and Heath 2008). Our setting also minimizes collective learning and group-level variety-seeking (e.g., Ariely and Levav 2000) as every user maintains his own virtual house. Finally, our ability to track and record individual choices frees us from potential self-reporting biases on conformity (e.g., Jetten, Hornsey and Adarves-Yorno 2006) and non-robust inferences on individual adoption thresholds from collective outcomes (e.g., Granovetter 1978). Given the exclusion of these confounding factors, conformity in our experiment can be safely attributed to a subject's need to belong to his group of friends, and non-conformity to his need for distinctiveness among his friends.

3. Data and Analysis

3.1 Data

We have 13,301 subjects in the main dataset, with 6,586 in condition A and 6,715 in condition B. Table 1 presents summary statistics for all the variables used in our analysis.¹⁴ The dependent variable *Converge* is 1 if a subject decides to adopt the supposedly most popular color shown in the experimental message. By design, the popular color is always different from the subject's original choice. Therefore, converging to that popular color is an act of conformity. The *Adoption Rate* ranges from 50 to 100 as expected, with a mean of 74.75 across all subjects in our final dataset.

As subjects may vary in their tendency to conform, we construct demographic variables based on information that is self-reported when users register with Kaixin. The average *Age* of the subjects is 32. Consistent with prior studies of social-network users, females are more represented than males in this game: the mean of the dummy variable *Female* is 0.68.¹⁵ We also collect information on a subject's hometown ("Jia Xiang" in Chinese), which is reported by about 83.6 percent of the subjects in our dataset. For each hometown, we obtain data on its GDP per capita (*GDP per Capita*) from the *2007 China Statistical Yearbook*, published by the National Bureau of Statistics of China, to gauge its relative affluence.¹⁶ We also collect information on whether the subject lives in or outside of China. The dummy variable, *Outside China*, takes the value of 1 if

¹⁴ We check whether the subjects are properly randomized by regressing a user's allocation onto all of the demographic and network-related variables described above. When Treatment takes the value of 1, the subject is in the treatment condition and when it takes the value of 0, the subject is in the control condition. Table A1 in the Appendix presents our randomization check results. None of the independent variables is significantly correlated with the assignment of the user's condition, which provides evidence of successful randomization.

¹⁵ See http://en.wikipedia.org/wiki/Gender_differences_in_social_network_service_use, accessed in March 2018.

¹⁶ The ranks of cities in China based on their GDP per capita remain mostly unchanged over time. As a robustness check, we use ranks of the cities instead of their actual GDP per capita and obtain similar results.

the subject lives outside China at the time of registration and 0 otherwise. Only two percent of the subjects in our dataset live outside of China.

Finally, we collect information on two network-related variables. The variable *Number of Friends* records how many friends each subject has upon entering our experiment. It varies substantially as one can see from Table 1, ranging from 11 (by our experimental design) to 1,024, the maximum number allowed by the site. The average subject in our experiment has about 57 friends on Kaixin. *Kaixin Duration* measures the number of days between a user’s registration and the time when the experiment is conducted. The mean of this variable is around 539, suggesting that the average subject has been registered on Kaixin for about 1.5 years by the time of our experiment. The range of this variable is also wide, from 0 days to 1,041 days, suggesting that some subjects in our experiment are newly registered while others have been registered for almost three years.

3.2 Model-Free Evidence

To present model-free evidence, we plot in Figure 5 the proportion of conforming subjects against the adoption rate of the popular color. The figure shows that the proportion of subjects who choose to conform to the most popular color among their friends is below 10% for all levels of adoption rates, which suggests that subjects are generally reluctant to repaint their walls once they have made the initial color choice. As one can see from the dashed line, which represents the control group, when subjects do not receive an experimental message, their likelihood of conforming to the (non-displayed) popular color does not vary with the (non-displayed) adoption rate ($p=0.805$).

On the other hand, as one can see from the solid line, the proportion of conforming subjects first decreases and then increases as the adoption rate of that color increases from 50% to 100%.

The proportion starts at around 7% when the adoption rate is close to 50%, gradually decreases to around 2% when the adoption rate is close to 80%, and then increases to around 9% when the adoption is close to unanimity. While the proportions are low on average, the changes appear to be significant.

The U-shaped result may be surprising at first sight, but a careful consideration of the trade-off between the two needs can help explain the pattern. In our experimental setting, the user's existing color choice is always designed to be different from the displayed popular color. When the adoption rate is only slightly above 50%, many friends may have chosen the same color as the focal user. The group of friends with the same color in this case may be too inclusive and not satisfy either of the user's needs very well: it does not give the user a good sense of belonging as he is still a minority, and it does not give him a good sense of uniqueness. As a result, the user may choose to conform to the majority of his friends to gain a better sense of belonging. As the adoption rate increases, the distinctiveness of the group of friends with the existing color choice increases. Remaining within this group better serves the user's need for uniqueness, and he is more likely to stay without changing his color choice. When the adoption rate approaches 1, the group of friends with the existing color choice becomes extremely distinctive and the user may again feel too uncomfortable to stay in the group. The U-shaped pattern is consistent with Brewer's (1991) optimal distinctive theory, which suggests that people are most likely to identify with groups that are neither too inclusive nor too exclusive.

Overall, the average likelihood of conformity for the treated group is always higher than that of the control group, and the contrast between the U-shaped curve for the treated group and the relatively flat curve for the control group suggests that subjects are indeed processing the information presented in our experimental messages.

3.3 Regression Analysis

3.3.1 U-Shaped Conformity

We use regressions to confirm the patterns observed above and further explore how conformity changes with popularity of the choice and user characteristics. The dependent variable in our regression is *Converge*.¹⁷ To illustrate how conformity changes with the adoption rate, we create dummies for different ranges of the continuous adoption-rate variable. In particular, we round each adoption rate to the nearest multiple of 5% and then use a separate dummy variable for adoption rates near each of such multiples, ranging from 50% to 100%. We have a total of 10 dummy variables as a result and our empirical specification is

$$Converge_i = \sum_{j=1}^{10} \alpha_j AR_{ij} + \sum_{k=1}^{10} \beta_k T_i AR_{ik} + \Theta X_i + \Phi T_i X_i,$$

where i is the index for the user, AR_{ij} is the dummy variable which indicates whether user i is shown an adoption rate in range j , T_i indicates whether the user is treated, and X_i includes user i 's number of friends, age, whether the user is female, the duration the user has been on Kaixin, GDP per capita of the user's hometown, and whether the user currently lives outside China. We use linear probability models to ease the interpretation of interaction variables, and later use other specifications to test for robustness. In our analysis, 100% of the predicted probabilities lie between 0 and 1. As shown in Horrace and Oaxaca (2006) and Angrist and Pischke (2008), linear probability models with robustness standard errors in this case yield unbiased and consistent estimates.

Table 2 reports the main results of our regression analysis. The first column includes observations only from the treated group and shows that conformity is clearly U-shaped with the

¹⁷ Results from other models are consistent with those from linear models, as can be seen from Table 8.

adoption rate. As one can see from Column 1, convergence is lowest when the adoption rate is near 80%, averaging 2.2% (0.070-0.048), and increases as the adoption rate moves away in either direction. It averages around 7% at the 50% adoption rate and 9.2% (0.070+0.022) at the 100% adoption rate. The regression confirms the U-shaped pattern identified earlier in Figure 5.

In the next three columns of Table 2, we compare the patterns of convergence between the control and treatment groups. In Column 2, we interact the treatment with each range of adoption rates and find the following results. First, the treatment overall leads to a higher overall level of conformity.¹⁸ That is, when users learn about a recently popular color among their friends, they are more likely to make the same choice. This confirms our prior belief that while online interactions tend to be more loosely organized than offline activities, there may still exist significant normative influences within the users' local friend networks as, after all, the main purpose of such networks is to serve people's need to socialize with others.

Second, the convergence of a treated user, when compared with that of an untreated user, first decreases and then increases with that color's adoption rate. This can be seen from the significant coefficients of the interactions of the treatment and the dummy variables for each of the adoption rate ranges. The U-shaped pattern is once again confirmed and as before, the adoption rates that are associated with the lowest levels of convergence are close to 80%, suggesting that people are the most content when identifying with a relatively small group within their own social networks.

3.3.2 Conformity and User Characteristics

¹⁸ Admittedly, this could be due to a potential "saliency" effect which suggests that users may be more likely to choose a color after seeing it in the experimental message. We rule out this possibility in Section 4.2.

Finally, we discuss what user characteristics are associated with stronger or weaker conformity. Column 3 of Table 2 suggests that when we look across both the treated and untreated subjects, users who are female, registered earlier on Kaixin, or born in more affluent towns are on average less likely than others to conform, while users who currently live outside China are more likely to conform. Since control-group participants do not receive information on the popular color, their adoptions of the randomly generated popular color should be random, and these patterns are therefore likely to be driven by the treated users.

We confirm this conjecture by interacting user characteristics with the treatment in Column 4, the last column of Table 2. The results confirm that users who are female, registered earlier on Kaixin, or born in more affluent towns are on average less likely to conform upon learning the popular color, while those who currently live outside China are more likely to conform. In addition, we find that users who have a smaller number of friends or are older are also more likely to conform upon learning the popular color. Given that the majority of Kaixin users are young, female, and live inside China, these results suggest that minority users (e.g., male, older, outside China), users of weaker social-economic status (e.g., born in less affluent towns), and new users (e.g., more recent registration, fewer friends) on the site are more likely to conform when a norm becomes salient amongst their friends.

It is important to note that for a particular user, having a minority status on the entire platform does not always translate into having a minority status in his friend circle. For example, friends of a newly-registered user born in a non-affluent town may all be similar to himself. Our findings suggest that such a user is still more likely than the average user to conform to the majority choice within his friend circle, possibly due to a stronger need to belong to his local network given his overall minority status on Kaixin.

The outside-China result may seem counter-intuitive on first sight as China has traditionally been seen as a collectivistic culture and those who managed to go abroad should have become more individualistic over time. It has to be, however, interpreted against the particular experimental setting of our study: these out-of-country users chose to remain active on the site and that is how they became part of our experiment. While they have earned membership in their friend circles on this platform, they are still a small minority compared to other users. The fact that they remained active suggests that they have a high willingness to maintain their relationships with friends in China, leading to an overall stronger likelihood to conform.

Overall, our findings are consistent with theories in social psychology (e.g., Goffman 1963, Aronson 1972, Phillips and Zuckerman 2001) that highlight that the marginalized and those of lower status may have a stronger need to conform. When explaining the middle-status conformity theory, Phillips and Zuckerman (2001) highlight that “unless the lowest ranked candidates are observable outsiders, a simple negative relationship between status and conformity should be observed.” In our context, since friendship requires mutual agreement on Kaixin, there are no observable outsiders and we similarly find that network members with a weaker status have a stronger tendency to conform, although measures of status, as one can see from our results, depend heavily on the particular social network.

Since prior studies (e.g., Obukhova, Zuckerman and Zhang 2011, 2014) suggest that older generations in China are more likely to be conformist due to the Cultural Revolution and Mao’s influence, we test for such potential cohort effects. Specifically, we break our subjects into five age cohorts: subjects born after 1990, between 1980 and 1989, between 1970 and 1979, between 1960 and 1969, and between 1950 and 1959. As one can see from Table 3, subjects who are born between 1950 and 1959 exhibit much stronger conformity toward the most popular color among

their friends. Since the Cultural Revolution took place from 1966 until 1976, it was then that subjects born between 1950 and 1959 experienced their most formative years. As prior studies suggest, they may exhibit a higher need for conformity. On the other hand, the other age cohorts are not significantly different from the subjects born after 1990: As pointed out by Zou and Cai (2016), there have been important changes in the Chinese population in the past two to three decades regarding individualism and taste for popularity, possibly due to a more relaxed political atmosphere.

4. Extensions and Robustness

In this section, we present two extensions of our main analysis. The first one examines popular choices with an adoption rate that is lower than 50%, and the second one examines popular color choices among all Kaixin users and consider users exposed to information on such choices as an alternative control group. Our results appear to be robust in both extensions.

4.1 Adoption Rates Lower than 50%

We have so far focused on adoption rates over 50 percent so that the popular color is a majority choice among the user's friends. Since our goal is to understand the nuances of normative social influences, it is natural to focus on a majority choice. Nonetheless, we reserve an additional pool of users to check the effects of lower adoption rates. The experimental message shown to these users is the same as the message shown to the treated subjects in Condition B above, except that the adoption rate is lower than 50%. For example, a user could see: "Recently, the most popular color among your friends is Blue (23%)."

Table 4 presents the summary statistics of the larger pool of users when we include all possible adoption rates.¹⁹ As there are six different colors for the subjects to choose from, the adoption rates in this larger sample range from 16 percent to 100 percent. As one can see from the table, the summary statistics for all the variables are similar except for the adoption rate, which is by definition lower for the larger pool.

Table 5 reports the results from the same set of regressions as before, using this larger sample. As one can see from the table, as well as Figure 6, our main results still hold. First, the treatment continues to have a significant impact on the likelihood of adoption, conforming in the presence of social influence. Second, when the adoption rate is higher than 50%, conformity remains U-shaped for the adoption rate of the popular choice. Third, as before, users who are female, with a higher number of friends, registered earlier on Kaixin, or born in more affluent towns are less likely to conform, while users who are older and living outside China are more likely to conform.

The first two columns of the table also suggest that a user's likelihood of adopting the popular color generally increases with the adoption rate before it hits 50%: the coefficients of the dummies for the adoption rate ranges, as well as their interactions with the treatment, exhibit an overall increasing pattern as the adoption rate increases. On the other hand, when we introduce user characteristics in the last two columns, this upward trend disappears while the U-shaped pattern for higher adoption rates continues to hold. It is possible that lower adoption rates of the most popular color are less informative on how popular the user's own color choice is. Overall, our results are more robust when the adoption rate of the popular choice is higher than 50%.

¹⁹ Table A2 in the Appendix presents the randomization checks and none of the explanatory variables is significantly correlated with treatment.

4.2 Popular Choices on the Entire Kaixin Platform

While our experimental design has carefully minimized many confounding factors, we could have introduced saliency or warning effects into the data. In particular, subjects could switch to or avoid a particular color after seeing it mentioned in our message simply because the message has made that color more salient to them. Alternately, they could switch to the popular color to avoid seeing the same message again. To ensure that the treated users are processing the information contained in our experimental message, rather than simply acting out of the saliency and warning effects, we reserve a third group of users in Condition C for further testing.

Users in Condition C are exposed to an experimental message that is similar to that in Condition B, except that the randomly generated color is now claimed to be popular on the entire Kaixin platform. The adoption rate, as before, is randomly generated and displayed in the same sentence as the popular color. As an example, a user in Condition C could see the following message: “Recently, the most popular color on the Kaixin platform is Blue (63%).” We use Condition C as an alternative control group to test for saliency and warning effects. If users are acting mainly out of such effects, we expect their reactions to the popular color on Kaixin to be similar to their reactions to the popular color among their own friends.

Table 6 presents the summary statistics with this alternative control group and confirms that user characteristics in this extension (Conditions B and C) are similar to those of the main analysis above (Conditions A and B).²⁰ Figure 7 presents model-free evidence on how users react to these two types of popular color as the adoption rate increases. As one can see, while the likelihood of adopting the popular color among one’s friends seems to vary significantly with the

²⁰ Table A3 in the Appendix presents randomization checks and further confirms that the two groups, B and C, do not significantly differ along any important dimensions.

adoption rate, the likelihood of the popular color on the entire platform does not. The contrast is especially sharp when the adoption rate exceeds 50%, with the first likelihood higher than the second on average.

We confirm these patterns through regression analysis, the results of which are presented in Table 7. As expected, we find that treatment has a significant effect, suggesting that the users' likelihood of adopting the most popular color on Kaixin is on average lower than their likelihood of adopting the most popular color among their friends. While users may converge to the popular color on Kaixin out of the saliency and warning effects, or a simple desire to follow the latest trend on the platform, their reactions to friends' choices are motivated by additional or stronger social influences. The result suggests that the treated subjects are indeed reading our experimental messages. Their stronger reaction to friends' choices makes intuitive sense as, after all, a user's choices can only be seen by his friends and not by the general public on Kaixin.

The regression results in Table 7 confirm again that a user's likelihood of adopting the majority choice among his friends changes significantly, in a U-shaped fashion, with the adoption rate of that choice. Since the control group is now different, the new result suggests that users pay careful attention to their friends' color choices since they see each other's houses, and they may not care as much about the proportion of all Kaixin users who have recently adopted a certain color. In other words, the relevance of the larger comparison group (all Kaixin users) may be too low for the users to care about the popularity of a particular choice.

Consistent with our previous results, users who are female and born in more affluent towns are less likely than other users to adopt the most popular color when the reference group changes from the entire Kaixin platform to the users' own friends. On the other hand, users who live outside of China are more likely to adopt the most popular color.

5. Concluding Remarks

While a large literature on social influences has documented conformity under normative influences, few studies have examined the impact of the popularity of a choice on conformity. Our large-scale field experiment on a leading Chinese social networking site suggests that conformity to a popular neutral choice among a user's friends may first decrease and then increase with the adoption rate of that choice. The U-shaped pattern is a new finding that is enabled by our ability to randomize popular colors and the associated adoption rates at the individual level. The new finding is consistent with existing theories in social psychology, such as optimal distinctiveness theory, which suggests that users would prefer to identify with social groups that are neither too inclusive nor too exclusive.

We also find that users who are male, older, born in less affluent towns, living outside China, registered later on Kaixin, or with fewer friends on the platform, are more likely to conform to the most popular color among their friends upon being given the information. Since the majority of users on Kaixin are female and young, our findings suggest that users of minority, newer, or lower social-economic status may be more likely to conform. These findings are consistent with prior theories (e.g., Goffman 1963, Aronson 1972, Phillips and Zuckerman 2001) that suggest that the users of a weaker status on a social network may have a stronger need to earn the approval of other group members by displaying conforming behavior.

From a methodological point of view, our research contributes to the literature by proposing an experimental design that removes confounding factors such as homophily, learning, and identity signaling to non-friend members. After removing these contextual factors, we can examine more fundamental needs such as the need to belong and the need for individuality. Our

findings provide useful insights into the design and optimization of normative marketing campaigns that have witnessed mixed levels of success so far (e.g., Schultz et al. 2007).

To maximize the neutrality of the choice, our study focuses on a descriptive norm. Existing campaigns, on the other hand, could feature injunctive norms, different cultural backgrounds and consumption choices that are not readily observable to others such as the in-home use of water and electricity. To gain a deeper understanding of the effectiveness of normative campaigns, future studies can dig deeper into, for example, how to select the reference groups to induce the greatest conformity.

References

- Abrams D (2009). Social identity on a national scale: Optimal distinctiveness and young people's self-expression through musical preference. *Group Processes and Intergroup Relations* 12(3):303–317.
- Angrist JD, Pischke J-S (2008) *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Ariely D, Levav J (2000) Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of Consumer Research* 27(3):279–290.
- Aronson E (1972) *The Social Animal* (Viking Press, New York).
- Asch SE (1952) *Social Psychology* (Prentice-Hall, Englewood Cliffs, NJ.)
- Austen-Smith DJr, Fryer RG (2005) An economic analysis of “acting white.” *Quarterly Journal of Economics* 120(2):551–583.
- Badea C, Jetten J, Czukur G, Askevis-Leherpeux F (2010). The bases of identification: When optimal distinctiveness needs face social identity threat. *British Journal of Social Psychology* 49(1): 21–41.
- Banerjee AV (1992) A simple model of herd behavior. *Quarterly Journal of Economics* 107(3):797–817.
- Bearman PS, Brückner H (2001) Promising the future: Virginity pledges and first intercourse. *American Journal of Sociology* 106(4):859–912.
- Benjamin, DJ, Choi JJ, Strickland AJ (2010) Social identity and preferences. *American Economic Review* 100(4):1913–1928.
- Berger J, Heath C (2008) Who drives divergence? Identity signaling, outgroup dissimilarity, and the abandonment of cultural tastes. *Journal of Personality and Social Psychology* 95(3):593–607.
- Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5):992–1026.
- Brewer MB (1991) The social self: On being the same and different at the same time. *Personality and Social Psychology Bulletin* 17(5):475–482.
- Brewer, MB (1993) Social identity, distinctiveness, and in-group homogeneity. *Social Cognition* 11(1): 150–164.

- Campo S, Cameron KA (2006) Differential effects of exposure to social norms campaigns: A cause for concern. *Health Communication* 19(3):209–219.
- Charness G, Rigotti L, Rustichini A (2007) Individual behavior and group membership. *American Economic Review* 97(4):1340–1352.
- Chartrand TL, Bargh JA. 1999. The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology* 76(6):893–910.
- Chen Y, Li SX (2009) Group identity and social preferences. *American Economic Review* 99(1):431–457.
- Cialdini RB (2003) Crafting normative messages to protect the environment. *Current Directions in Psychological Science* 12(4):105–109.
- Cialdini RB, Reno RR, Kallgren CA (1990) A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology* 58(6):1015–1026.
- Codol J-P (1981) Une approche cognitive du sentiment d'identité [A cognitive approach to the feeling of identity]. *Social Science Information / Sur les sciences sociales* 20(1):111–136.
- Currarini S, Jackson MO, Pin P (2009) An economic model of friendship: Homophily, minorities, and segregation. *Econometrica* 77(4):1003–1045.
- Deutsch M, Gerard HB (1955) A study of normative and informational social influences and individual judgment. *Journal of Abnormal and Social Psychology* 51(3):629–636.
- Gerber AS, Green DP, Larimer CW (2008) Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review* 102(1):33–48.
- Goette L, Huffman D, Meier S (2006) The impact of group membership on co-operation and norm enforcement: Evidence using random assignment to real social groups. *American Economic Review* 96(2):212–216.
- Goffman E (1963) *Stigma: Notes on the Management of Spoiled Identity* (Prentice-Hall, Upper Saddle River, NJ).

- Granovetter M (1978) Threshold models of collective behavior. *The American Journal of Sociology* 83(6): 1420–1443.
- Griskevicius V, Goldstein NJ, Mortensen CR, Cialdini RB, Kenrick, DT (2006) Going along versus going alone: When fundamental motives facilitate strategic (non)conformity. *Journal of Personality and Social Psychology* 91(2):281–294.
- Haines MP, Barker GP, Rice R (2003) Using social norms to reduce alcohol and tobacco use in two Midwestern high schools. *The Social Norms Approach to Preventing School and College Age Substance Abuse* (Jossey-Bass, San Francisco).
- Hartmann WR, Manchanda P, Nair H, Bothner M, Dodds P, Godes G, Hosanagar K, Tucker C (2008) Modeling social interactions: Identification, empirical methods, and policy implications. *Marketing Letters* 19(3):287–304.
- Hofstede G (1980) *Culture's Consequences* (Sage, Beverly Hills, CA).
- Hornik R, Jacobsohn L, Orwin R, Piesse A, Kalton G (2008) Effects of the national youth anti-drug media campaign on youths. *American Journal of Public Health* 98(12):2229–2236.
- Hornsey MJ, Jetten J (2004) The individual within the group: Balancing the need to belong with the need to be different. *Personality and Social Psychology Review* 8(3):248–264.
- Horrace WC, Oaxaca RL (2006) Results on the bias and inconsistency of ordinary least squares for the linear probability model. *Economics Letters* 90(3):321–327.
- Jetten J, Hornsey MJ, Adarves-Yorno I (2006) When group members admit to being conformist: The role of relative intragroup status in conformity self-reports. *Personality and Social Psychology Bulletin* 32(2):162-173.
- Kallgren CA, Reno RR, Cialdini RB (2000) A focus theory of normative conduct: When norms do and do not affect behavior. *Personality and Social Psychology Bulletin* 26(8):1002–1012.
- Kim H, Markus HR (1999) Deviance or uniqueness, harmony or conformity? A cultural analysis. *Journal of Personality and Social Psychology* 77(4):785–800.

- Kossinets G, Watts DJ (2009) Origins of homophily in an evolving social network. *American Journal of Sociology* 115(2):405–450.
- Lau RR (1989) Individual and contextual influences on group identification. *Social Psychology Quarterly* 52(3):220–231.
- Leonardelli GJ, Loyd DL (2016) Optimal distinctiveness signals membership trust. *Personality and Social Psychology Bulletin* 42(7):843-854.
- Leonardelli GJ, Pickett CL, Brewer MB (2010) Optimal distinctiveness theory: Framework for social identity, social cognition, and intergroup relations. *Advances in Experimental Social Psychology* 43(2): 63-113.
- Lieberson S, Mikelson KS (1995) Distinctive African American names: An experimental, historical, and linguistic analysis of innovation. *American Sociological Review* 60(6):928–46.
- Manski CF (1993) Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60(3):531–542.
- Markus HR, Kitayama S (1994) A collective fear of the collective: Implications for selves and theories of selves. *Personality and Social Psychology Bulletin* 20(5):568–579.
- Mas A, Moretti E (2009) Peers at work. *American Economic Review* 99(1):112–145.
- Obukhova E, Zuckerman EW, Zhang J (2011) Why was there no fashion in Mao’s China? Effects of politics on culture in the case of Chinese naming practices. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1878678.
- Obukhova E, Zuckerman EW, Zhang J (2014) When politics froze fashion: The effect of the Cultural Revolution on naming in Beijing. *American Journal of Sociology* 120(2):555–583.
- Perskins HW, Linkenbach JW, Lewis MA, Neighbors C (2010) Effectiveness of social norms media marketing in reducing drinking and driving: A statewide campaign. *Addictive Behaviors* 35(10):866-874.
- Perkins HW, Craig DW, Perkins JM (2011) Using social norms to reduce bullying: A research intervention among adolescents in five middle schools. *Group Processes & Intergroup Relations* 14(5):703–722.

- Phillips DJ, Zuckerman, EW (2001) Middle status conformity: Theoretical restatement and empirical demonstration in two markets. *American Journal of Sociology* 107(2):379–429.
- Scharfstein DS, Stein JC (1990) Herd behavior and investment. *American Economic Review* 80(3):465–479.
- Schultz PW, Nolan JM, Cialdini RB, Goldstein NJ, Griskevicius V (2007) The constructive, destructive, and reconstructive power of social norms. *Psychological Science* 18(5):429–434.
- Schultz PW (2010) Changing behavior with normative feedback interventions: A field experiment on curbside recycling. *Basic and Applied Social Psychology* 21(1):25–36.
- Snyder CR, Fromkin HL (1980) *Uniqueness: The Human Pursuit of Difference* (Plenum, New York).
- Tajfel H, Billig MG, Bundy RP, Flament C (1971) Social categorization and intergroup behaviour. *European Journal of Social Psychology* 1(2):149–178.
- Triandis HC (1995) *Individualism and Collectivism* (Westview, Boulder, CO).
- Tucker C, Zhang J (2011) How does popularity information affect choices? A field experiment. *Management Science* 57(5):828-842.
- Vignoles VL, Chryssochoou X, Breakwell GM (2000) The distinctiveness principle: Identity, meaning, and the bounds of cultural relativity. *Personality and Social Psychology Review* 4(4):337–354.
- Zhang J, Zuckerman EW, Obukhova E (2016) A lack of security or of cultural capital? Acculturative conservatism in the naming choices of early 20th century US Jews. *Social Forces* 94(4):1509–1538.
- Zou X, Cai H (2016) Charting China's rising individualism in names, songs, and attitudes. *Harvard Business Review*. <https://hbr.org/2016/03/charting-chinas-rising-individualism-in-names-songs-and-attitudes..>

Figure 1: Screenshot of the “Virtual Homes” Application



Figure 2: Screenshot of the Message Window for a Random Subject in Condition B



Figure 3: Design of the Experiment

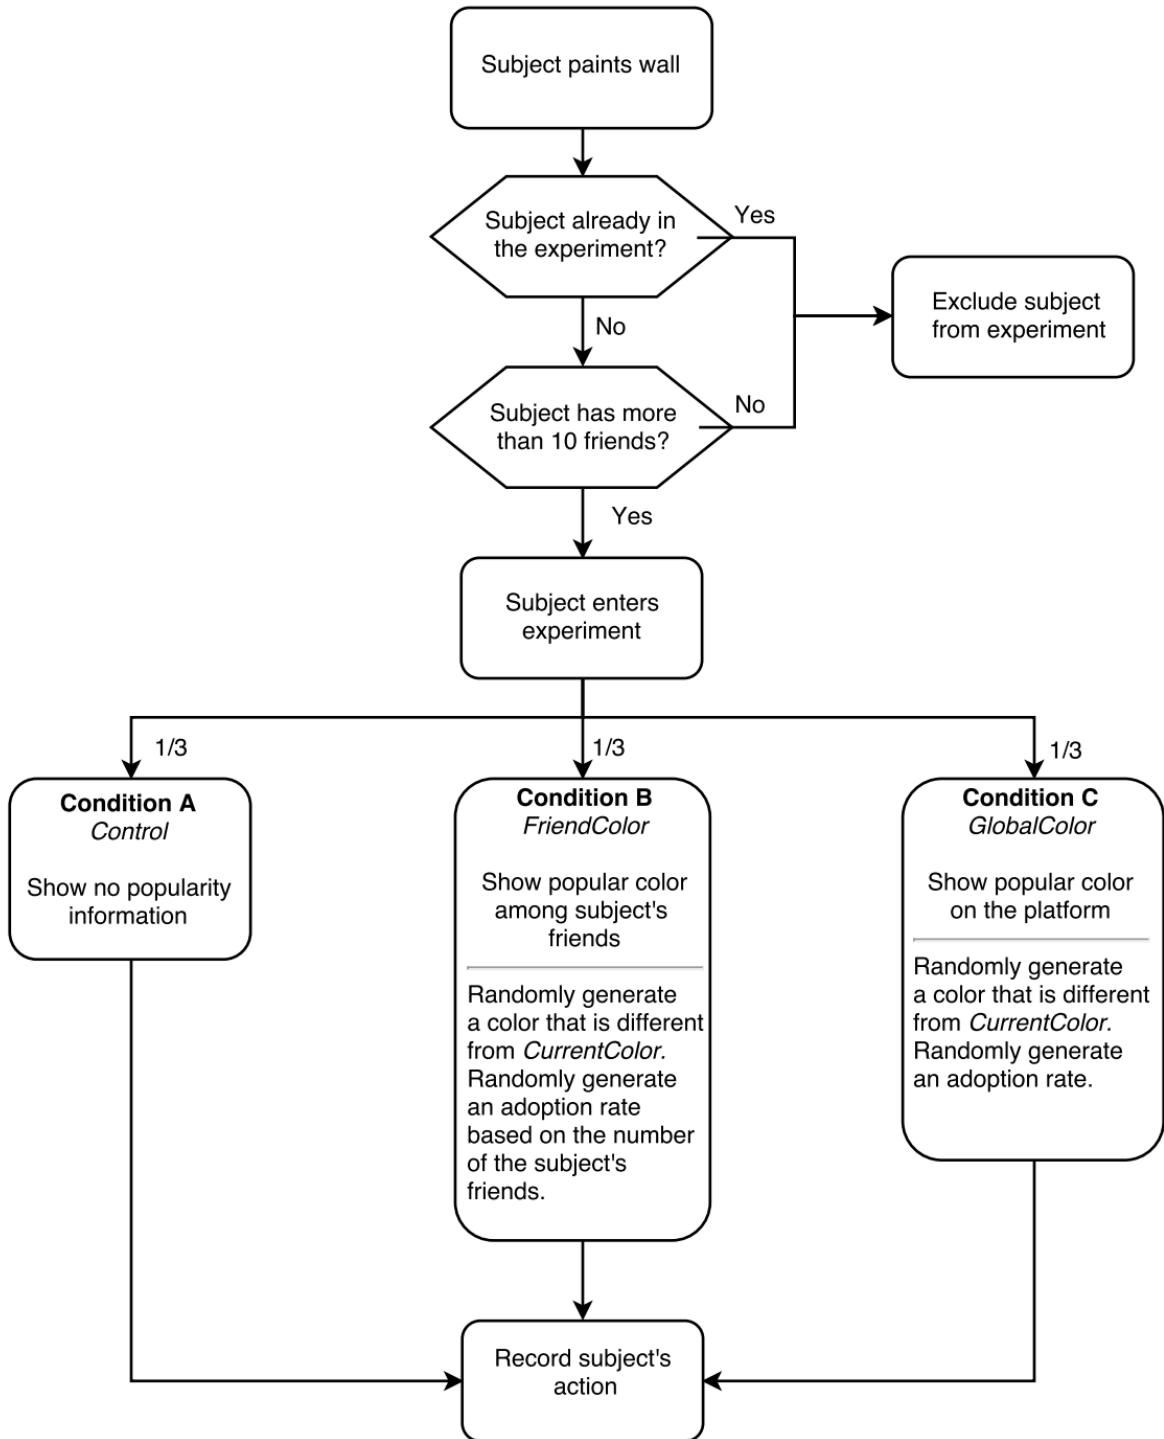
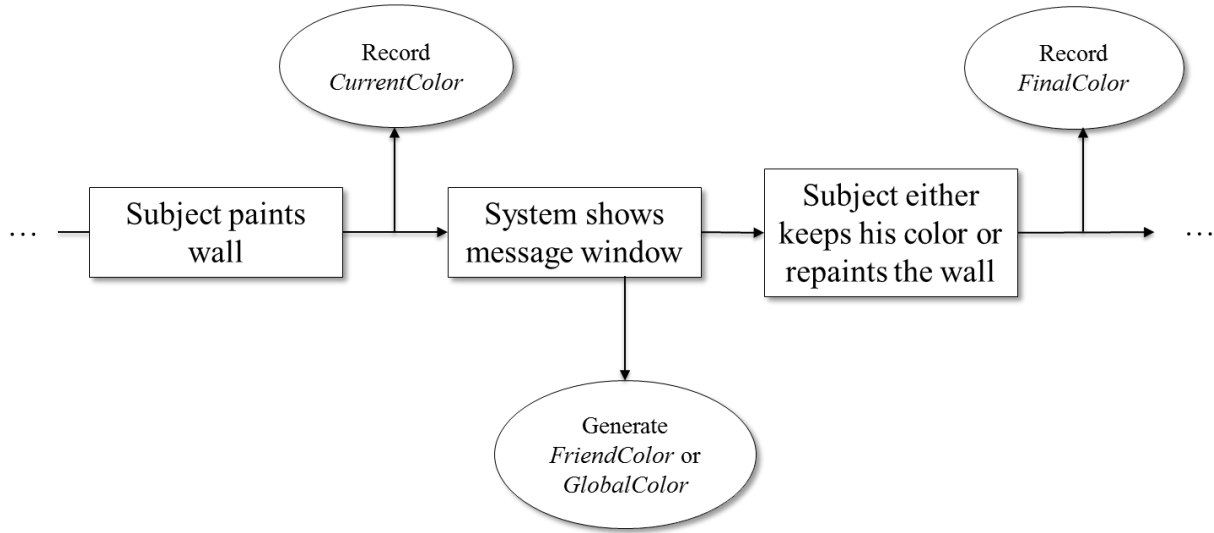



Figure 4: Timeline of Events and Color Records




Note: The *before* choice made before the subject sees the experimental message is recorded as *CurrentColor* and the final choice made after he sees the message is recorded as *FinalColor*. The randomly generated most popular color among a subject's friends is recorded as *FriendColor* and that among all Kaixin users is recorded as *GlobalColor*. The circles denote the points in time at which messages are generated and colors are recorded.

Figure 5: U-Shaped Conformity towards Majority Choice among Friends




Note: For each subject, Converge is 1 if the subject chooses the popular color and 0 otherwise. We round all adoption rates to the nearest multiple of 5% and then compute the average of Converge for subjects within each adoption rate bracket. Error bars indicate 95% confidence intervals.

Figure 6: Conformity Patterns with Full Range of Adoption Rates



Note: For each subject, Converge is 1 if the subject chooses the popular color and 0 otherwise. We round all adoption rates to the nearest multiple of 5% and then compute the average of Converge for subjects within each adoption rate bracket. Error bars indicate 95% confidence intervals.

Figure 7: Conformity Patterns with Global Color as Control



Note: For each subject, Converge is 1 if the subject chooses the popular color and 0 otherwise. We round all adoption rates to the nearest multiple of 5% and then compute the average of Converge for subjects within each adoption rate bracket. Error bars indicate 95% confidence intervals.