Short on shots:
Can calls on self-restraint be effective in managing the scarcity of flu vaccines, and what do they reveal about behavior[†]

Alain de Janvry, Elisabeth Sadoulet, and Sofia Villas-Boas

University of California at Berkeley

July 16, 2007

Abstract

We designed an experiment at the time of the 2004 flu vaccine shortage providing information about the sharply reduced number of clinics and an appeal on self-restraint to a campus population. This information induced a net increase in vaccines distributed, and perversely the net increase originated entirely in non-priority individuals. It is those who had used vaccination services the previous year that increased most their demand, demonstrated least self-restraint, and contributed most to cheating. Analysis of their stated reasons for wanting a vaccine suggest that this is more likely due to loss aversion than to risk aversion. The surprising finding is that it is the priority population that exercised self-restraint.

JEL Classifications: C93. Keywords: Randomized experiment, shortage, self-restraint, cheating.

[†] Address: University of California at Berkeley, 207 Giannini Hall, Berkeley, CA 94720-3310; e–mails: alain@are.berkeley.edu, sadoulet@are.berkeley.edu, sberto@are.berkeley.edu.

Salience

"We knew that once people heard there was a shortage, more people would try to get the vaccine." *San Francisco Chronicle*, October 11, 2004

Self-restraint

"There is a strong spirit of cooperation during this crisis. We have no intention of taking any draconian steps to enforce this state of emergency." *San Francisco Chronicle*, October 9, 2004

Cheating

"Flu shots, often a test of bravery, became a test of character ..., and not everyone was passing." *San Francisco Chronicle,* October 7, 2004

I. INTRODUCTION

While history is replete with situations where societies have been confronted with unexpected commodity shortages, the way shortages have been managed has been quite varied. When a market exists, rising prices serve as the main rationing device, with targeted subsidies eventually used to ease the burden of adjustment on designated groups considered at risk. When the price is fixed, allocation of the scarce commodity across wanting individuals has to be done by introducing rules to distribute the commodity to those presumed most in need. These rules can be implemented by screening and/or by calls on self-restraint. However, information about the shortage, including to motivate the call on self-restraint, also induces an increase in demand associated with greater salience of the commodity and increased eagerness in acquiring it. Self-restraint by some in refraining from acquiring the commodity can thus be countervailed by increases in demand by others.

Given these contradictory behavioral responses, the net effect may be in favor of self-restraint, leading to an aggregate decline in demand and good targeting, or of increased demand and/or poor targeting. While, in the long run, initiatives can be taken to respond to the shortage by increasing supply, understanding what motivates the short-run demand responses to the shortage is important to help better manage scarcity in a non-market setting. In particular, policy makers would like to know how effective can broad-scale calls on self-restraint be in managing the shortage of a vital good since this is likely to be a less politically costly approach than coercive screening. This is the first objective that we address in this paper. The second objective is to use the observed behavioral responses to identify which categories of individuals responded to information about scarcity and to calls on self-restraint, and to reveal the apparent motivations behind these responses.

To fulfill these objectives, we set up a randomized experiment to decompose responses to the large unexpected flu vaccine shortage that occurred in the Fall of 2004. Because the approach followed by health authorities was to manage the shortage by a call on voluntary restraints, we use the observed behavioral responses to the experiment to measure how far calls on cooperation can go in managing scarcity. The experiment took place at a flu clinic held at a California university campus medical center. Prior to the clinic, we subjected the campus population to two randomized experimental treatments: in treatment one (T_1), a group of departments received an email informing about the reduced number of vaccination clinics (scarcity) and their corresponding schedule (deadlines); in treatment two (T_2), another group of departments received an email with the same information as T_1, but additionally appealing (as the Center for Disease Control was recommending at the time) for non-members of defined priority groups to refrain from seeking vaccination. The rest of the campus population did not receive an email from us and served as a control group C. Two weeks after this clinic, the medical center sent an email to the campus population announcing a last clinic.

This randomized design, and the surveys done at the two clinics, allow us to analyze both the demand for vaccination and the actual distribution of vaccines. Demand was measured by the population that came to the clinic seeking vaccination. Actual distribution was to those who did not walk away when informed of screening and who were not rejected by the clinic soft-screening. To analyze demand, we decompose quantitatively the different behavioral responses at play: the response to information about scarcity and deadlines is measured by the difference in behavior between T_1 and C; the response to calls on self-restraint, conditional on information about scarcity and deadlines, by the difference in behavior between T_2 and T_1; and the net effect of these two types of responses by the difference in behavior between T_2 and C. The relative contribution of subgroups in the campus population to each type of response can also be identified. Results show a very large effect of information on scarcity and deadlines in increasing demand, particularly among non-priority people, which was only partially counteracted by voluntary restraints. Among non-priority people, most of the demand originated in users of vaccination the previous year, showing what can be interpreted as strong loss aversion for a commodity they had learned to value. Priority groups, by contrast, responded equally strongly to scarcity as to calls on self-restraint, resulting in a wash on demand, even though the calls on self-restraint were not directed at them. To analyze the actual distribution of vaccines, we decompose the roles of information, self-restraint, and screening. An analysis of confidential self-declared membership in priority groups and of unusual patterns of declared membership in priority groups provides evidence on the extent of cheating among candidates for a flu vaccine. In the end, the strategy failed to reserve the scarce resources to the targeted population. The number of vaccinations distributed increased by 17% and all the addition went to non-priority people.

II. BEHAVIORAL RESPONSES TO A SHORTAGE: LESSONS FROM THE LITERATURE

2.1. Increased salience as a response to scarcity

It is well recognized that perceptions of scarcity can induce a sharp increase in demand due to rising salience of the scarce good, worsening whatever true shortage there might be. Some of the great famines in history like those in Bengal in 1943, Ethiopia in 1973, and Bangladesh in 1974 in fact occurred without any disruption in supply (Sen, 1981). The "Great Toilet Paper Shortage" caused in zest by Johnny Carson in 1973 also occurred without any change in supply.[1] In other cases, the scarcity effects of shortfalls in supply were greatly amplified by induced consumer buying. In a market setting, given a contraction in supply, if demand expands in response to the shortage, then the price increase is greater than the one caused solely by the leftward supply shift. With fixed prices, the "panic buying" effect induced by a fall in supply is amplified by lack of price response, requiring some type of rationing device. Examples are the oil "buyer panics" of 1971 and 1973 that resulted in long lines at the gas pumps as government froze prices, with time waiting in line becoming the rationing device (Adelman, 2004). That scarcity enhances desirability has long been recognized in the marketing literature (Folger, 1992; Lynn, 1992a and 1992b). The 2004 flu vaccine shortage analyzed here was similarly managed under price control.[2] A rise in demand was fully expected to happen as a response to the shortage, and rules were introduced to direct scarce supplies toward priority groups.[3] Because the commodity is of vital importance for people at risk, information about scarcity to justify a call on

[1] In his Late Night Show monologue, Johnny Carson said: "You know what's disappearing from the supermarket shelves? Toilet paper. There's an acute shortage of toilet paper in the United States." The consequence of this statement made in the early 1970's, a time of shortages -- oil in particular --, was that the next morning many of the 20 million television viewers ran to the supermarket and bought all the toilet paper they could find. By noon, most of the stores were out of stock since, despite trying to ration it, they couldn't keep up with demand.

[2] The few cases of price gouging resulted in legal charges.

[3] As a county Public Health Department spokesperson said: "We knew that once people heard there was a shortage, more people would try to get the vaccine." *San Francisco Chronicle*, October 11, 2004.

self-restraint among people not in priority groups also creates greater salience of vaccination among non-priority individuals, resulting in an obvious dilemma for the management of scarcity without strict screening.

2.2. Decreased procrastination as a response to scarcity

Another response that can increase demand as a consequence of a shortage is that the strict deadlines associated with rationing in the distribution of a scarce good that will eventually run out may reduce the occurrence of normal-time procrastination. Procrastinators are individuals who delay tasks until a later period, and who, when the later period arrives, delay those tasks again and again if there are no strict deadlines for getting things done (Akerlof, 1991). Sirois et al. (2003) found empirical evidence that procrastination applies as well to decisions related to individuals' own health. Procrastination can be overcome by the introduction of strict deadlines. This is consistent with studies that find, for example, that if manufacturers place a deadline on redemption of the coupons they distribute, the probability of redemption increases (Silk, 2004); and that the shorter the time students are given to complete a task, the lesser the likelihood that they will fail to complete it (Tversky and Shafir, 1992). If procrastinators postpone getting a flu shot in normal times when there are no deadlines, even among individuals in priority groups, strict deadlines introduced by the rationing scheme may induce many of them to overcome delaying and seek vaccination, adding to the rise in demand induced by the shortage.

2.3. Self-restraint as a response to scarcity

Voluntary restraints in response to a call on cooperation can be expected to hold when there is clear information about expected benefits, effective monitoring and enforcement, and repeated interactions. For this reason, this is more likely to occur in small groups with long time horizons (Olson, 1965). In this perspective, responses to broad-scale demands for voluntary restraints in the face of scarcity can be expected to fail. Yet, there is also abundant evidence of willingness to cooperate in situations of relatively anonymous and sporadic relations. A number of recent behavioral experiments (e.g., Fehr and Gächter, 2000; Gintis et al., 2003) have found that individuals behave more cooperatively than the "self-interest individual model" would predict (Rabin, 1998). This applies, for instance, to tax payment where the observed rate of tax abidance cannot be explained by current levels of audit risks and penalties (Feld and Frey, 2002). "Tax morale" needs to be invoked to explain observed levels of compliance. Willingness to cooperate is possible even in large social groups as it can be motivated by the desire for social approval (Holländer, 1990), by conforming to social norms for fear that non-compliance by oneself will lead to their collapse (Azar, 2004), or by satisfaction in cooperating if it helps improve one's self-image (Trivers, 1971).

In calling on broad-scale cooperative self-restraint to manage a flu vaccine shortage, the expectation was that individuals not in priority groups would voluntarily incur the risk of being sick to allow the scarce resource to reach the people most in need. On a campus where many people know each others and identify with others, one could further expect restraint to be motivated by a sense of community and a concern with reputation.

2.4. Loss aversion as a response to scarcity

Introducing calls on self-restraint among non-priority individuals implies that previous users of vaccination services will agree to drop use of the service, and also that previous non-users will refrain from becoming first-timers. Research in psychological economics by Kahneman and Tversky (1979) and Kahneman et al. (1991) showed that a person's well-being depends not only on his current consumption of goods but also on how current consumption compares to past consumption. Prospect theory suggests that there is a strong asymmetry in the evaluation of an increase and a decrease in consumption (Tversky and Kahneman, 1991). People care much more about losses relative to their reference point than about gains (Bowman, Minehart, and Rabin, 1999). As a consequence, we would expect that previous non-priority users of vaccination services would be much more reticent to respond to calls on self-restraint than individuals who did not previously use vaccination. In the analysis of responses to the flu vaccine crisis, it will consequently be important to contrast the behavior of previous users motivated by loss aversion to that of previous non-users to whom this logic does not apply.

III. EXPERIMENTAL DESIGN AND DATA COLLECTION

3.1. The flu vaccine shortage and the timeline of events

On Monday, October 4, 2004, the campus medical center in our experiment sent its routine annual reminder that everyone should receive a flu shot every year and informing of the schedule for the six planned vaccination clinics starting with October 6 and ending in December, 2004. On Tuesday, October 5, half of the U.S. supply of flu vaccine was pulled back from the market because of possible contamination.[4] Starting on Wednesday, October 6, numerous media articles about the flu vaccine shortage started to inform the American public. The United States Center for Disease Control (CDC) appealed to the public for people not in defined priority groups to voluntarily forego vaccination. On October 6, the campus medical center held the first of its six previously scheduled vaccination clinics. Two days later, on Friday October 8, it announced on its website reduction to only two in the number of subsequent clinics, with occurrence of the originally announced other three subject to vaccine availability. On Saturday, October 9, some California counties declared an emergency to enforce a State directive restricting vaccination to priority groups. The county where the campus is located did not at that time officially announce enforcement of this directive.[5]

On Monday, October 11, one week after the shortage was first announced, the two experimental treatment emails (T_1 and T_2) were sent out to the campus population. Monday the 11th was a national holiday and on the next day, Tuesday October 12, the second clinic, henceforth referred to as clinic A, took place, offering flu shots to the campus population and the non-campus community, and soft-screening candidates. This screening measure was not previously announced by the medical center. Individuals had to sign an affidavit declaring that they belonged to one of the priority groups, but with no proof asked as health-center personnel were more concerned with servicing than with policing. Individuals did not even have to specify which priority group they belonged to. These groups were: children 6-23 months of age, adults 65 years of age and older, women expecting to be pregnant during the flu season, health care workers with direct patient care, out-of-home care givers, individuals with household contacts of children less than 6 months old, adolescents on chronic aspirin therapy, and persons ages 2 through 64 with a chronic medical condition (such as asthma, diabetes, heart disease, chronic kidney disease, or who had chemotherapy or immune-compromised conditions). These groups had always been designed as priority, even in previous years when vaccination was recommended to all. We conducted our first survey during clinic A.

On Wednesday the 13th, the campus medical center cancelled all remaining clinics and recommended the population to check for updates. The update came two weeks later. On

[4] British regulators cut the U.S. vaccine supply in half by condemning 48 million doses at a Liverpool factory owned by Chiron Corporation, a U. S. company based in Emeryville, California, after bacterial contamination was found.

[5] "There is a strong spirit of cooperation during this crisis," said the respective County Public Health Officer. "We have no intention of taking any draconian steps to enforce this state of emergency." *San Francisco Chronicle*, October 9, 2004.

Wednesday, October 27, the medical center sent a campus-wide email informing about the date for a final clinic and announcing that, given the shortage, all candidates for flu vaccination would be asked to sign an affidavit that they belong to one of the priority groups. By the time of this last clinic, that we henceforth call clinic B, screening of participants was common practice across the U.S. and, most likely, the information sent via email to the campus population was by then also known to the non-campus community. Signature of an affidavit was required from all candidates, certifying membership in one of the priority groups. However, no hard proof of qualification into one of these groups was requested by the screening personnel. On Monday, November 1, we conducted our second and last survey during clinic B.

3.2. The experiment at clinic A

We randomly selected departments to receive two different kinds of email treatments. Members of the first subset of departments (T_1) received an email informing that only two clinics would be offered (scarcity) and giving the dates for these clinics (deadlines). Members of the second subset (T_2) received an email containing the same information as sent to T_1 plus, in accordance with CDC recommendations at that time, a call on self-restraint in seeking vaccination for people not in priority groups.[6] The priority groups were described in detail in the T_2 email. The remaining departments (the control group C) received no email.

The experimental design was thus intended to allow identification of the following behavioral responses:

- From comparison of the T_1 and C groups, the impact on vaccine demand and distribution of information about scarcity and deadlines.
- From comparison of the T_2 and T_1 groups, the impact on demand and distribution of sending a call on self-restraint, conditional on information about scarcity and deadlines.
- From comparison of the T_2 and C groups, the net impact on demand and distribution of sending information about scarcity and deadlines, and of calling on self-restraint.

Emails were sent to faculty, staff, and graduate students by the management services officers (MSO) of the different departments. Of the 69 departments on campus, 10 were drawn for each of the treatment groups. However, 3 that had been selected for the second treatment did not follow up immediately upon our email. Given the extremely tight schedule of the experiment, this prevented them from participating to the experiment, leaving therefore 10 departments for T_1,

[6] Self-restraint is here defined here as "being informed of the reduced number of clinics, and not coming to a clinic in response to the call for the population not in priority groups to defer vaccination".

7 for T_2, and the remaining 52 for *C*. The emails to undergraduate students were sent by the student affairs officers (SAO) for declared majors corresponding to the selected departments and by the dean of the college for undeclared students. With 3 selected departments not having undergraduates and 6 SAO not responding immediately, the experiment included 8 majors and the undeclared from one college for T_1 and 3 majors and the undeclared from one college for T_2, leaving the rest for *C*. The numbers of treated faculty, staff, and students in the T_1, T_2, and *C* groups are given in Table 1. We verify below that the randomization process is valid by performing tests of difference in characteristics across treatment groups for each of the faculty, staff, graduate, and undergraduate samples. Of the campus population of 38,604, 8,695 were in T_1, 12,233 in T_2, and 17,676 in *C*.

As the opportunity of getting a vaccine was offered at the workplace, it is likely that social interactions among co-workers influenced individual decisions to go to the clinic. This can be due to the transmission of information received, to mutual influence in appreciating the value of getting or not getting vaccinated for the flu, or to the fact that people who work together may go together to the clinic, a fact that we observed at the clinics. These social interactions take place regardless of any treatment effect, including in the control groups. They, however, also affect the treatment effect itself, in so far as the treatment of one person has spillovers on the other members of the social network. Our experiment is not set up to distinguish the direct influence of the email treatment from the indirect influence that would occur within a department, as all members of the same professional category in a department received the same information.[7]

The validity of our analysis in measuring the effect of sending an email relies on the stability assumption, i.e., that there was no interference across treatment units. Although this is not a guarantee that social interactions did not affect the experiment, clinic A occurred the day after a national holiday, giving people limited time to interact on the morning of October 12, the day of clinic A, after they potentially read their emails. By sending the emails through administrative channels, we also believe that it minimized the chances of social interactions across departments.[8]

[7] This is in contrast with the experimental design used in Duflo and Saez (2003) who subjected a random sample of members of a subset of departments in a campus population to treatments. Their objective was to assess the role of information and peer-effects on the decision to enroll in employer-sponsored Tax Deferred Account retirement plans. They found the interesting result that when treating only 50% of the department members, the indirect effect through department co-workers is almost as high as the direct effect of the treatment.

[8] There is little motivation to forward the email to people outside the department since each email recipient believes it is likely that members of other departments were also receiving such email directly from their

3.3. The data and randomization tests

For clinic A, no screening had been announced. Yet, the list of qualifying priority groups was posted at the entrance of the medical center, and some screening was performed by the registration personnel. Among candidates for flu vaccination, some walked away upon reading the list of priority groups, others were screened out by the center personnel.

The survey forms were filled out by basically everybody. This may be either because the survey looked official, or because the opportunity cost of completing the survey while waiting on line was very small. We also surveyed the people who came in and, upon seeing the poster and noticing the screening, decided on their own to forgo vaccination. Information collected includes age, gender, campus affiliation by department and professional category, whether individuals got a flu shot in each of the last three years, and the reasons for them to come which included membership to the different priority groups.

On the day of clinic B, the response rate was again close to 100% once we started handing out the survey forms.[9] The survey questionnaire was extended to ask whether individuals were or not in each of the priority groups in 2003.

At clinic A, 738 individuals filled questionnaires, with 427 from campus and 311 from the non-campus community.[10] Out of the 394 campus members with departmental information, 37% were from the treatment group T_1, 25% from the treatment group T_2, and 38% from the control group. At clinic B, 610 persons filled questionnaires, 385 from campus and 225 from the non-campus community.

Because the randomization was based on departments and because departments have different configurations in terms of faculty, staff, and student composition (see Table 1), each of these groups gives rise to stratified samples of the campus population. Estimates of campus population statistics are thus derived from weighted averages of statistics by professional category.

---

official administrative units. When asked how they had heard about the clinic, and in particular if it was through an email received from campus during the last two days, 17% of participants from *C* said they did. Despite this response, we are quite confident that very few members of *C* have effectively received this email. This is because there may have been some confusion with the campus-wide email sent only one week earlier. In addition, the question itself may have induced a "yes" response to something that may have appeared as if it should have been true.

[9] This time, the clinic started about one hour earlier than announced to accommodate the long lines, so our survey team missed the first hour of people who came to the clinic.

[10] This includes staff from the campus administrative units as they were not part of the experimental design.

To verify the validity of the experimental design, we perform randomization tests on observables for the three groups T_1, T_2, and C. Relevant dimensions that could affect behavior toward vaccination on which we have information are gender, age in 10 categories, race, occupation, and wage category, for faculty and staff; and gender for graduate and undergraduate students. Results are reported in Table 2 for pair-wise differences between T_1, T_2, and C. The similitude between the three randomized groups is excellent for all four groups. Only two comparisons (faculty occupation in T_2 vs. T_1 and staff age in T_2 vs. C) fail the randomization test at the 10% significance level among 36 pairs, a failure rate of only 5.6%. Comparisons between T_1, T_2, and C can thus be taken with confidence. In section 4.2, we confirm this by running robustness checks on average treatment effects by including age and gender as control variables in the estimated equations.

IV. IMPACT OF TREATMENTS ON DEMAND FOR A FLU VACCINE

From the total number of members of each professional category for each department or major, we compute participation rates to the clinics for each unit of randomization. Note that, as the information was collected at the clinics themselves, we do not have individual data on the characteristics of the population of interest, namely the whole campus population. For reasons of privacy, we only obtained from the campus administration the distribution of some characteristics by department and professional category, not individual data. To analyze differences in participation rates by subsets of the campus population along characteristics that we observe in the survey but not in the population at large (whether the individual had been vaccinated in 2003 or not and whether the individual is member of a priority group or not), we thus have to rely on the assumption of orthogonality of the treatment to all such characteristics in each professional category.

4.1. Impact analysis: Average treatment effects

We are interested in measuring the impacts of the treatments on demand for a flu shot. Demand under treatment T is defined as the proportion of the campus population that would come to the clinic under that treatment. Equivalently, it is the average probability that a random campus member will come to clinic A under treatment T:

$$P\left(Y^A = 1 | T\right),$$

where Y^A is a binary variable equal to 1 if an individual comes for a flu shot at clinic A, and *T* is the treatment (*C*, T_1, or T_2). Obviously, for each individual member of campus, we only observe one of the three potential outcomes $Y^A|T$. To estimate campus demand under the different treatments, we rely on the randomization scheme and compute campus participation as a weighted average of the participation rates by professional category *K* and department *D* in each treatment group:

$$\hat{P}\left(Y^A = 1|T,K\right) = \sum_{D\in T} \hat{P}\left(Y^A = 1|T,K,D\right)\frac{N_{KD}}{\sum_{D\in T} N_{KD}}$$

$$\hat{P}\left(Y^A = 1|T\right) = \sum_{K} \hat{P}\left(Y^A = 1|T,K\right)\frac{N_K}{N}$$

where $\hat{P}$ are observed participation rates, N_{KD} and N_K are population of category *K* in department *D* and on campus, respectively, and *N* is total campus population. Variances of these estimators are weighted averages of the binomial variances for each professional category and department. Standard errors thus include clustering effects that could be due to social interactions among members of the same professional category in each department.

The demand $\hat{P}\left(Y^A = 1|T\right)$ under each treatment $T \in \{C, T_1, T_2\}$ and the average treatment effects $\hat{P}\left(Y^A = 1|T\right) - \hat{P}\left(Y^A = 1|T'\right)$ of *T* relative to T' in clinic A are reported in Table 3 for the four professional categories and for the whole campus population. Results show that sending information about the reduced number of clinics ($T_1 - C$) induced a doubling (110% increase) of the demand for flu vaccine (from 0.8 to 1.6% of the campus population) due to a large increase across all professional categories. The difference between the two treatment groups ($T_2 - T_1$) measures the effect of sending a call on self-restraint in addition to information on scarcity and deadlines. This effect on behavior was to decrease demand to 1.0% of the campus population, or by 37.5%, with the largest effects among undergraduate students (-50.9%) and faculty (-42%). These two effects resulted for the whole campus population in a significant 31.2% net increase between *C* and T_2. This net increase was largely due to staff and undergraduate students. While this shows that calls on self-restraint can indeed be heard, reduced demand by cooperators was far from sufficient to compensate for the increase in demand induced by information about scarcity and deadlines. Analyzing the heterogeneity of responses to the information treatments in section 4.4 will help interpret the source of the increase in demand and show that it is not solely due to the re-scheduling of planned attendance to clinics that were cancelled.

4.2. Robustness checks

Of the demographic variables that are available for the campus population at large, gender and age are also available in the survey, allowing robustness check on the measure of average treatment impacts, controlling for these two observables using regression analysis. For each of the four campus population categories, OLS estimation of the average treatment effects on demand for a flu vaccine, first with no controls, and then adding age or gender as control variables, are reported in Table 4. The estimates without any control are identical to the direct estimates reported in Table 3. Results are remarkably stable to the addition of age and gender covariates, confirming that the randomization was successful in making treatment orthogonal to the other factors influencing demand for vaccine.

4.3. Intention to treat vs. treatment

The treatment effects of *sending emails* that we have thus far analyzed can be considered as either interesting in themselves, to the extent that they measure the impact of a well defined type of information campaign, or because they provide measures of the intention to treat effects of the treatment defined as *being informed*. In the latter case, the sending of emails can be viewed as an encouragement to be informed. The intention to treat is usually interpreted as a conservative estimate of the treatment effect, although it need not be the case if there is a direct effect of the email beyond the information that it conveys (Hirano, Imbens, Rubin, and Zhou, 2000).

Following the literature on encouragement design, one can distinguish four subpopulations. The never-takers are individuals who remain untreated despite the encouragement. In our experiment, they are the (never-informed) individuals that have been sent an email but have not read it, and remained uninformed. The subpopulation of always-takers consists of the (always-informed) individuals who are informed regardless of the encouragement. These people found on their own the information on the reduced number of clinics and their schedule or have heard of the Center for Disease Control's call on self-restraint from other sources. Given the very short time between the decision to restrict the number of clinics and the clinic itself and the lack of any direct publicity by the health center, this information probably did not spread much (cancellation of the last three clinics was decided and posted on the medical center's website on October 8 but not further advertised). However, the call on self-restraint from the Center for Disease Control may have been heard more widely. There is no subpopulation of defiers in this experiment.[11] The last subpopulation, the compliers, consists of the individuals who acquired the information from the email sent to them.

Under the condition that there is no direct effect of the email on the always-informed and the never-informed, sending the email only has an effect on the compliers. The effect of sending an email (intention to treat effect) is thus a conservative measure of the effect of the *being informed* treatment.

That the never-informed are not influenced by the encouragement design is not controversial. However, for the always-informed, receiving the email provides a reminder effect in addition to the information effect. In this case, the reminder effect of the October 11 email may not have been large since a campus-wide reminder had been sent only a week before. We, however, cannot assess its importance in this context.

How does the effect of "*sending an email*" relate to the effect of "*receiving (or having read) the email*"? Denote by *RT* the treatments "have read the information on the reduced number of clinics and their schedule" and "have read the call on self-restraint from the CDC, in addition to the information on the remaining clinics". These two *RT* treatments differ from the treatments that we analyze as not everyone has read its email that same day.[12] However, because of the random design of the email treatments, we can assume that having read one's email is orthogonal to the treatments *T*, and obtain intention to treat effects of *RT* that are simple scaling downs of the treatment effects by the email reading rate. This can be seen as follows.

The conditional probability of coming to clinic A can be decomposed into:

$$P\left(Y^A = 1 | T\right) = P\left(Y^A = 1 | Read = 1, T\right) P\left(Read = 1 | T\right) + P\left(Y^A = 1 | Read = 0, T\right) P\left(Read = 0 | T\right).$$

With randomized treatment, the percentage of the campus population that has read its email that same morning is independent of the treatment, i.e., $P\left(Read | T\right) = P\left(Read\right)$. In addition, for the population that has not read the email, treatment is de facto ineffective, i.e., $P\left(Y^A = 1 | Read = 0, T\right) = P\left(Y^A = 1 | Read = 0\right)$. Therefore:

$$P\left(Y^A = 1 | T\right) = P\left(Y^A = 1 | Read = 1, T\right) P\left(Read = 1\right) + P\left(Y^A = 1 | Read = 0\right) P\left(Read = 0\right).$$

Treatment effects are thus scaled down by the probability that the emails were read:

[11] These would be individuals who would be informed without the email but not informed when sent an email.

[12] 75% of the persons that came to the clinic from T_1 and 69% of those from T_2 said that they had learned about the clinic schedule from an email received in the last two days.

$$P\left(Y^A = 1 | T\right) - P\left(Y^A = 1 | T'\right) = \left[P\left(Y^A = 1 | RT\right) - P\left(Y^A = 1 | RT'\right)\right] P\left(Read = 1\right).$$

In summary, our experiment allows us to strictly measure the effect of sending emails (providing a reminder effect on opportunity for vaccination together with new information on scarcity and the reduced number of clinics, and the call on self-restraint from the Center for Disease Control) to all the members of a department. The measured effect includes the social network effect internal to the department. It is proportional to the effect of receiving an email, with the scaling factor equal to the email reading rate. Its relationship to the effect of the information itself is mitigated by several issues that could make it either higher or lower than the pure effect of the information contained in the emails.

4.4. Insights on the roles of information and calls on self-restraint in influencing the demand from specific groups

The effects of information, calls on self-restraint, and combined information and self-restraint on the demand for vaccination from specific groups are presented in Table 5. We contrast the behavioral responses of members and non-members of priority groups and, within these categories, of previous users (people who were vaccinated in 2003) and first-timers (people who were not vaccinated in 2003). The reason to contrast the responses of members and non-members of priority groups is that the scarcity management strategy was expected to induce an increase in demand from the first due to information and a decline in demand from the latter due to self-restraint. The reason to contrast the responses of previous users to that of first-timers is that they can occur through different rationales:

i) Rescheduling: Part of the increase in demand due to information on scarcity and deadlines originates with people who had planned to participate in a clinic that was cancelled and who rescheduled their visit to the health center to one of the two available clinics. Rescheduling is done by both previous users and first-timers.

ii) Increased salience: Demand for vaccination increases due to greater appreciation of the importance of getting vaccinated conveyed by information about scarcity. This is a genuinely new demand from people who would otherwise not have wanted to be vaccinated. Hence, it specifically applies to the response of first-timers.

iii) Reduced procrastination: Procrastinators are individuals who wanted to be vaccinated in the past (have a positive latent demand), but always postponed doing so under normal times as there were no strict deadlines.[13] Imposition of strict deadlines as a consequence of scarcity prevents them from postponing vaccination forever, resulting in an increase in effective demand. This behavioral response also applies specifically to first-timers.

iv) Loss aversion: Previous users will reschedule in response to information and may resist calls on self-restraint due to loss aversion, possibly augmented by salience. With the experimental set up we use, we are not able to distinguish between the specific effects of loss aversion and of some fundamental unobserved heterogeneity between previous users and first-timers, such as risk aversion. However, in the context of a scarcity crisis, loss aversion may be a good explanation of observed behavior. We can disentangle some of these effects by taking advantage of having stated reasons for seeking flu shots in the survey questionnaire. In particular, we are able to distinguish the reasons given by previous users and first-timers in seeking flu shots. If previous users state less risk related reasons compared to first-timers, then we can reasonably attribute their coming to loss aversion rather than to risk aversion.

To determine what we can measure with the available data, let X = 1/0 indicate if the individual was vaccinated in 2003 (previous user). X is not observed in the population at large. Hence, we cannot estimate the participation rate $P\left(Y^A = 1 \middle| X, T\right)$ conditional on X. What we can estimate is the probability of being a participant with characteristic X:

$$P\left(Y^A = 1, X \middle| T\right)$$

or the relative impact of T_1 on the probability of participation conditional on X,

$$\frac{P\left(Y^A = 1 \middle| X, T = T_1\right)}{P\left(Y^A = 1 \middle| X, T = C\right)} = \frac{P\left(Y^A = 1, X \middle| T = T_1\right) \big/ P\left(X \middle| T = T_1\right)}{P\left(Y^A = 1, X \middle| T = C\right) \big/ P\left(X \middle| T = C\right)} = \frac{P\left(Y^A = 1, X \middle| T = T_1\right)}{P\left(Y^A = 1, X \middle| T = C\right)},$$

where $P\left(X \middle| T = T_1\right) = P\left(X \middle| T = C\right)$ because of orthogonality of T_1 and C to X. Similar expressions can be written for the other treatment comparisons, as well as for the priority/non-priority groups.

In Table 5, there are four notable results in the heterogeneity of responses:

i) There was a much larger increase in demand due to information (T_1 – C) among non-members of priority groups (219.9%) than among members of priority groups (55.8%). Increase in salience due to information about scarcity thus mobilized a huge response among those not in priority to receive a vaccine.

[13] Although there was always a limited number of clinics offered in previous years at this particular medical center, vaccination remained available at many other places in the area throughout the season.

ii) There was a moderate response to calls on self-restraint across all groups, reducing demand in the range of 30 to 39%.

iii) Among non-members of priority groups, most of the increase in demand originated with previous users, largely attributable to rescheduling, a behavior that (as we shall see in section 4.5 and Table 7) can be associated with loss aversion, more than risk aversion.

iv) Members of priority groups demonstrated remarkable self-restraint that cancelled their increase in demand due to information about scarcity and deadlines. That self-restraint would result in no increase in demand among this group was clearly not part of the scarcity management strategy. The consequence is that the totality of the increase in demand originated with non-members of priority groups, a definitely unintended effect as well.

v) To investigate whether responses to appeals on self-restraint by members of priority groups could be related to reputation effects, we take advantage of having asked in the questionnaire whether an individual was present on the Berkeley campus during previous years. The conjecture would be that if you are new on campus you are less identified with the community and less worried to be recognized at the clinic by people who may think that you should not be there, and therefore would be less responsive to the appeal on self-restraint. Indeed, we find that there is a large difference in the point estimates of change in demand (even if not statistically significant) between those who have been on campus for at least one year (-41.8) and the newly arrived (-17.7) (result not reported in Table 5). The result thus supports the interpretation that self-restraint was related to a greater sense of community and/or a greater concern with reputation.

4.5. Who were the first-timers?

The exceptional rise in first-timer demand in 2004 versus 2003 is also suggestive of the role of information on demand. Denote by Y_t the indicator for having received a flu vaccine in year *t* (2002 or 2003). The proportion of first-timers (meaning not having been vaccinated the previous year) in the population that was vaccinated in year *t* is the conditional probability $P\left(Y_{t-1}=0 \middle| Y_t = 1\right)$. We do not observe this ratio in the population at large, although we do observe it in the population that came to get a flu shot at clinic A, $P\left(Y_{t-1}=0 \middle| Y_t = 1, Y^A = 1\right)$.

Using standard conditional probability relationships, we can write:

$$P\left(Y_{t-1}=0 \middle| Y_t = 1\right) = P\left(Y_{t-1}=0 \middle| Y^A = 1, Y_t = 1\right) \frac{P\left(Y^A = 1 \middle| Y_t = 1\right)}{P\left(Y^A = 1 \middle| Y_{t-1} = 0, Y_t = 1\right)}.$$

We make, in addition, the reasonable assumption that the probability of coming to clinic A, conditional on having received a flu vaccination in year *t*, is independent of whether one had or not received a flu vaccination the previous year *t* - 1:

$$P\left(Y^A = 1 \middle| Y_{t-1} = 0, Y_t = 1\right) = P\left(Y^A = 1 \middle| Y_{t-1} = 1, Y_t = 1\right) = P\left(Y^A = 1 \middle| Y_t = 1\right).$$

This gives: $P\left(Y_{t-1} = 0 \middle| Y_t = 1\right) = P\left(Y_{t-1} = 0 \middle| Y^A = 1, Y_t = 1\right)$, meaning that the observed ratio of first-timers each year in the population that came to clinic A measures the share of first-timers in the population at large. Combining members and non-members of priority groups, Table 5 shows that 29% of the increase in demand due to information about scarcity observed in 2004 came from first-timers.

Results in Table 6 show a sharp increase in first-timers' demand for vaccination in 2004 compared to previous years.[14] This is seen by the incidence of first-timers for flu vaccination among those who came to clinic A this year, compared to the incidence of first-timers in the previous year, in the non-campus community and in campus group *C* in clinic A, and in the non-campus community in clinic B. These are the three groups that did not receive any special information from campus about deadlines or affidavits, and hence who were responding to general knowledge about scarcity. At clinic A, 11.9% of non-campus community participants were first-timers in 2004, compared to rates of first-timers of 3.1 and 3.7% the two previous years. In group *C*, 25.7% were first-timers in 2004, compared to rates of 13.5% and 5.9% the two previous years[15]. The phenomenon of rising demand was even sharper at clinic B, with information on shortage more widely available in the press. At this clinic, 22.6% of non-campus community participants were first-timers in 2004 compared to rates of 5.8% and 3% the two previous years.

These sharp increases in first-timers for flu vaccines could be due to any year 2004 effect. However, the dominant phenomenon that year was greater information in the media about the existence and importance of flu vaccination, and about the existence of a shortage. We can thus conclude that, as expected from the literature on responses to scarcity, the spread of information

[14] In 2004, demand is measured by "coming to the clinic to seek vaccination". For the previous years, we use "has received a flu vaccine" as demand since there was no restriction.

[15] Although six clinics were announced in 2002 for this health center, delays in shipment disturbed the announcement of clinic dates, which were progressively scheduled as vaccines became available, and at the end only five clinics were effectively held. To the extent that unreliable supply and uneven announcements discourage potential newcomers more than regular customers, this could explain a lower value for the ratio of first-timers in 2002 compared to 2003.

about a fall in supply led to a sharp increase in demand from people who had never requested a flu shot before.

Why did first-timers come to the clinics compared to previous users? To answer this question, Table 7 compares first-timers and previous users at clinic A in terms of membership in priority groups, and other reasons invoked for desiring a flu shot. We restrict the analysis to members from the community and campus groups C and T_1, which were not subject to the call on restraint. The results are quite revealing of who the first-timers are. While 64.2% of the previous users are members of a priority group, this applies to only 33.4% of the first-timers. The reasons these non-priority members invoke are that they cannot afford to miss a day of work or study, that they are concerned with a potential epidemic, and other reasons that relate to risk of contagion due to exposure to others (living in dorms, being in contact with people, traveling abroad). Hence, first-timers are driven by anxiety, salience, and decreased procrastination more than by seriousness of medical consequences, which would qualify them as priority.

4.6. Suggestive evidence for loss aversion and procrastination

The interpretation of loss aversion behavior among non-members of priority groups who had been previous users is supported by contrasting the reasons given to seek a flu shot between previous users and first-timers of the flu shot clinics among non-priority groups, reported in the bottom part of Table 7. By contrast to the first-timers who invoked shortages and epidemic, the reasons they give to seek a flu vaccine are clearly not related to risk.

At clinic B to which only members of priority groups came as screening had been announced, 2/3 of the first-timers are people who were in the same priority group last year, and yet had not been vaccinated (the other third being people who became at risk this year). This certainly includes a response to the increased salience of vaccination and in particular to the explicit focus this year on priority groups. However, with vaccination campaigns targeting every year these same priority groups, most were certainly aware of the importance for them of being vaccinated, suggesting that we are seeing among them many procrastinators responding to imposition of strict deadlines as a consequence of the shortage.

V. THE VACCINE RECIPIENTS: EVIDENCE ON CHEATING

5.1. Evidence on cheating under screening

How can cheaters be detected? The anonymous survey, filled by candidates for a flu shot, asked for a confidential self-declaration as to whether the person belonged or not to each priority group, with the possibility of belonging to more than one. Some people walked away after filling the questionnaire as they admitted not belonging to any priority group. For those who remained in line, the medical personnel engaged in soft verification (with no proofs asked) that the individual qualified for receiving a vaccination. Screening was unexpected at clinic A, but fully expected at clinic B as it was explicit in the clinic announcement. All candidates for a flu shot thus had to officially announce membership in one of the priority groups in order to be considered for vaccination, had they declared confidentially in the survey that they were in one or not. The screening nurse then decided to accept or reject the candidate. We thus have information from each candidate for vaccination about: (1) whether self-declared in a priority group or not, and (2) whether the individual received a flu shot or not (as he either walked away or was denied). This allows us to construct four categories of candidates in columns 1 through 4 of Table 8:

- Effective screening: These are the candidates who declared in the survey not belonging to a priority group and who were not serviced, either because they walked away by themselves or were screened out by the center staff. Many of them might have been uninformed about the call for self-restraint and screening (screening was not announced for clinic A), while others probably came with the intention to cheat (the schedule for clinic B was always given with information that screening would be enforced).
- Legitimate service: Those are the candidates who declared in the survey belonging to the priority groups and were indeed serviced.
- Exclusion error (Type II): Those are the candidates who declared belonging to the priority groups, but were however denied a flu shot. While this could be a genuine exclusion error, it is more likely a category of persons that were properly detected not being priority while they self-declared being priority in an attempt to cheat.
- Inclusion error (Type I): Those are non-priority persons who were serviced (cheaters). They probably spoke the truth in the survey, but still orally declared being in a priority group to the staff, signed the affidavit, and were not screened out.

Effective screening, revealing lack of information or intention to cheat, was unimportant for non-campus community participants (column 1): the rejection rate was very low (2.9% at clinic A and 1.9% at clinic B). However, this was not the case among campus candidates at clinic A where it reached 25.1% in group *C* and was significantly higher in T_1 (36.1%) and T_2 (32.5%) than in *C*. While non-priority candidates may have come to the clinics because of lack of information on the existence of priority groups, this could not be the case for at least campus group T_2 at clinic A and for the whole campus population at clinic B (where screening had been announced). And yet, it is interesting that screening was higher in the treatment group T_2 than in the control group, although again not precisely measured and not significantly higher because of the small number of observations. This suggests that attempting to cheat the system was reinforced by anxiety created by information about scarcity, even when accompanied by explicit calls on self-restraint.

Legitimate service (column 2) was almost universal in the non-campus community (92% at clinic A and 97.2% at clinic B). It was also high among campus participants to clinic B (88.2%). It was low, however, among campus participants to clinic A, and lower in the treatment groups T_1 and T_2 than in *C* due to the importance of screening and cheating for these participants.

Exclusion errors, whereby members of priority groups were denied vaccination, were almost non-existent in both clinics and for all groups (column 3). Screening by nurses was thus on the side of concern with exclusion errors, at the cost of greater inclusion errors. If the health center's objective was to weight exclusion errors more heavily than inclusion errors, to make sure that a minimum number of people at risk would be left un-serviced, then screening was indeed very effective.

Finally, in column 4, cheaters are those who self-declared not being in a priority group, yet were given a flu shot. Cheating behavior is better measured as the share of non-members of priority groups that obtained a vaccine as reported in column 6. This share is very high among non-campus community members, which is understandable given the higher cost for them of getting to the clinic. Among campus members, it increased from 24.1% in *C*, to 28.3% in T_1, and to 33.8% in T_2.

The contrast between first-timers and previous users is also quite revealing of who the first-timers are. First-timers contain a greater share of individuals uninformed and/or intent on cheating, both in the control and treatment groups. However, among non-priority people (column 6), it is previous users that accounted for the highest share of cheaters, motivated by what we interpreted as loss aversion.

5.2. Impact of the information campaign, call on self-restraint, and soft screening on distributed vaccines.

In Table 9, we report the impact of treatments T_1 and T_2 on the number of vaccines distributed, contrasting vaccinations given to members and non-members of priority groups. Sending information by email reminding people that there was a shortage and hence only two remaining clinics increased substantially not only the demand for vaccines (+110.1% in Table 3) but also the number of vaccines distributed after screening (+76.6% in Table 9). Calls on self-restraint induced a significant decline in demand (-37.5% in Table 3) and in vaccines distributed (-33.6% in Table 9). In the end, information, self-restraint, and screening resulted in an increase of 17.2% (Table 9) in vaccines distributed to the campus population.

What is striking in these results is that while information was effective in bringing to the clinic a large number of members of priority groups that the CDC certainly wanted to vaccinate (a 52.8% increase in column T_1 - C), it induced a far greater increase in vaccination among non-priority groups (a 276.5% increase, or almost five fold), despite a certain level of screening at the clinic itself. The self-restraint response further exacerbated this contrast, inducing a somewhat larger decrease (although statistically not significantly different) in the priority population than in the non priority population (-35.4% vs. -27.6%). Hence, considering together information, self-restraint, and screening, the overall increase in vaccination (17.2%) is solely due to an extraordinary increase in vaccination of non-priority people (+172.6%). Their share in the vaccinated population rose from 11% to 25%. This indicates that cheating was indeed extensive among those who were vaccinated.

5.3. Evidence on cheating from the survey questionnaire

How else can cheaters be detected? What we used above to identify cheaters was presumed truthful self-reporting in the survey of not being in a priority category, and yet making it through scrutiny of the medical personnel and receiving a flu vaccine. There can, however, be cases where self-reporting may not have been truthful in spite of guaranteed anonymity. In this case, cheaters are people who falsely declared themselves to be in a priority category in the survey, did this again on the required affidavit, and were not detected by medical personnel because providing hard proof of being in the category was not demanded. How can we know that self-reporting was not truthful? Only if there are obvious statistical irregularities in the risk categories invoked.

One such irregularity that reveals cheating is in the pattern of ages declared. Figure 1 representing the distribution of self-declared ages is striking in showing a peak at age 65, preceded by a dip with missing numbers between ages 60 and 64. The 65 years old group is two to three times larger than the average per age between 66 and 70. This is true for non-campus community as well as campus participants, so we pool all the data from the two clinics in analyzing this pattern. Existence of an abnormally high number of participants of age 65 is formally analyzed with the estimation of an age profile for participants.

Discontinuity at age 65 is due to two effects: one is the age eligibility criterion that would imply a discontinuity between ages 64 and 65, with more participation of 65 years old; the other is cheating on age where people younger than 65 declare themselves to be 65. The discontinuity that would reveal cheating must consequently be measured from above. To do this, we estimate the age profile of participants 66 years old and above only, and predict from above the participation at age 65. We explored different functional forms (3rd degree polynomials in age, 1/(1+ age)). The estimated curves are reported in Figure 1.

Cheating at 65 is measured by the difference between observed and predicted number of participants to the clinics. Predicted numbers of 65 years old are 30.7 (standard error of 3.2) with the 3rd degree polynomial and 34.6 (standard error of 1.7) when function of 1/(1+ age). The observed number of 77 is more than twice the predicted values, estimated with relatively high precision. This suggests widespread cheating on age. Because there was no verification of age, most of this cheating could go undetected. Estimation of "missing" 61-64 years old is not precise as the profile of candidates between 50 and 61 years old is not smooth. However, the corresponding estimate of the distance between observed numbers and predicted numbers for these four age groups gives a missing number of 75.3 persons (standard error of 23.3).

VI. CONCLUSION

The first objective of this paper was to analyze the effectiveness of the strategy used to manage the 2004 flu vaccines scarcity crisis based on responses observed on a U.S. university campus. The strategy mandated by the Center for Disease Control consisted in the definition of priority groups and in a call on self-restraint by the rest of the population, supported by soft-screening. The expected outcome of the strategy was a reduction in demand to accommodate the shortage, while servicing the priority groups.

Results from the controlled experiment we set up to decompose the effects of information on scarcity/deadlines and of calls on self-restraint show that the outcome was quite different from what was expected both in terms of the magnitude of the response and of who responded.

Analysis of demand shows that there was a very large effect of information on scarcity and approaching deadlines (+110%), but that the call on self-restraint was also valuable, reducing demand by 38% (i.e., canceling 72% of the increase in demand due to information). However, in the end, the strategy could not prevent a 31% net increase in demand at the clinic where we ran the experiment. Neither could soft-screening of candidates prevent a 17% increase in vaccines effectively distributed compared to no strategy. The most disturbing result is that the increase originated entirely in non-priority people (+173%) whose share of distributed vaccines increased from 11% to 25% as a consequence of the strategy.

In terms of shortage management, the main lesson learned from this experiment is thus that calls on self-restraint supported by soft-screening appear to be insufficient to manage the scarcity of such a vital good as a flu vaccine. This lesson in crisis management should be remembered as repeats of flu shortages may be looming in the future.

The second objective of the paper was to identify the behavioral responses of different types of individuals and the motivations behind their behavior. Disaggregating by types of individuals shows that the largest perverse response came from non-priority previous users of flu shots. The increase in demand due to information was by far the largest among them as it quadrupled. And response to the call on self-restraint was particularly modest among them, resulting in a large net increase in demand of 163%. In the end, it is the increased demand by this particular group that was the main driving force in making the strategy ineffective: their demand increased sharply in response to scarcity and deadlines, their self-restraint was modest, and they contributed most to cheating in going undetected through soft-screening.

In contrast to the behavior of non-priority previous users, the most unexpected result was the uncalled-for self-restraint observed among priority individuals that fully erased their desired increased share of distributed vaccines.

Analysis of the determination to cheat shows that it reinforced access to vaccines by non-priority individuals. Because there was soft-screening at the clinic, the possibility was offered of walking away when learning about the screening or of trying to get through and be vaccinated despite admitting of not being member of a priority group. Based on self-declaration of being member or not of a priority group, the percentage of cheaters among non-priority individuals increased from C, to T_1, and to T_2, and was higher among previous users than among new candidates (although measured imprecisely). This shows more determination of getting vaccinated by these same non-priority users whose demand responded more to information about scarcity and approaching deadlines, and less to the call on self-restraint. How cheating occurred is evident on age declaration, with the number of "65 years old" more than twice the predicted value, while about half of the predicted 61-64 years old are missing.

To reveal the motivations behind the observed behavioral responses to the shortage and the call on self-restraint, we take advantage of the field experiment and field survey that asked detailed questions on individual characteristics, recall on past characteristics, and also reasons for seeking flu shots. We found that while first-timers indicated clear concern with risk, previous users did not invoke risk as their motivation to seek vaccination. This suggests that the behavior of non-priority previous users can be interpreted as strong loss aversion to a service they had received before, rather than risk aversion.

Unexpected self-restraint among the priority population is correlated to length of stay on campus. We use this to infer that this behavior may be motivated by greater community identification and by reputation concerns.

We thus conclude that loss aversion, decreased procrastination, and cheating were all important behavioral determinants of increased demand responses to the shortage. Self-restraint among priority individuals, likely related to community identification and reputation effects, was an unexpected non-trivial behavioral response that decreased demand and helped ease the shortage.

REFERENCES

Adelman, M. A. 2004. "The Real Oil Problem." *Regulation* 27(1): 16-21.

Akerlof, G. A., 1991. "Procrastination and Obedience." *American Economic Review* 81(2): 1-19.

Azar, Ofer. 2004. "What Sustains Social Norms and How They Evolve? The case of Tipping". *Journal of Economic Behavior and Organization* 54(1): 49-64.

Bowman, David, Deborah Minehart, and Matthew Rabin. 1999. "Loss Aversion in a Consumption-Saving Model." *Journal of Economic Behavior and Organization* 38: 155-178.

Duflo, E., and E. Saez, 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *The Quarterly Journal of Economics* 118(3): 815-842.

Fehr, E., and S. Gächter, 2000. “Fairness and Retaliation: The Economics of Reciprocity.” *Journal of Economic Perspectives* 14(3): 159–81.

Feld, Lars, and Bruno Frey. 2002. “Trust Breeds Trust: How Taxpayers Are Treated.” *Economics of Governance* 3(2): 87-99.

Folger, Robert, "On Wanting What We do not Have, Commentary," *Basic and Applied Social Psychology*, 1992, 13 (1), pp.123-133.

Gintis, H., S. Bowles, R. Boyd, and E. Fehr, 2003. “Explaining Altruistic Behavior in Humans.” *Evolution and Human Behavior* 24: 153-172.

Hirano, Keisuke, Guido Imbens, Donald Rubin, and Xiao-Hua Zhou. 2000. “Assessing the effect of an influenza vaccine in an encouragement design.” *Biostatistics* 1(1): 69–88.

Holländer, H., 1990. “A Social Exchange Approach to Voluntary Cooperation.” *The American Economic Review* 80(5): 1157-1167.

Kahneman, D., and A. Tversky. 1979. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica* 47: 263-291.

Kahneman, D., J. Knetch, and R. Thaler. 1991. “Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias.” *Journal of Economic Perspectives* 5: 193-206.

Lynn, Michael, "The Psychology of Unavailability: Explaining Scarcity and Cost effects on Value," *Basic and Applied Social Psychology*, 1992a, 13 (1), pp.3-7.

Lynn, Michael, "Scarcity's enhancement of Desirability: The Role of Naive Economic Theories," *Basic and Applied Social Psychology*, 1992b, 13 (1), pp.67-78.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.

Rabin, Matthew. 1998. “Psychology and Economics.” *Journal of Economic Literature* 36(1): 11-46.

Sen, Amatya. 1981. *Poverty and Famines*. Oxford: Clarendon Press.

Silk, T. 2004. “Why Do We Buy But Fail to Redeem? Examining the Breakage in Mail-in-Rebates." Working paper, University of Florida.

Sirois, F.M., M. Melia-Gordon, and T. A. Pychyl. 2003. “ ‘I’ll Look After My Health Later’: An Investigation of Procrastination and Health.” *Personality and Individual Differences,* 35: 1167-1184.

Trivers, R. 1971. “The Evolution of Reciprocal Altruism.” *Quarterly Review of Biology* 46: 32-57.

Tversky, A., and D. Kahneman. 1991. “Loss Aversion in Riskless Choice: A Reference-Dependent Model.” *Quarterly Journal of Economics* 106: 1039-1061.

Tversky, A., and E. Shafir, 1992. “Choice Under Conflict: The Dynamics of Deferred Decisions." *Psychological Science* 3: 358-592.

Table 1. Number of faculty, staff, and students by random treatment and control groups

Campus population	C	T_1	T_2	Total
Professional categories				
Faculty	1,640	834	559	3,033
Staff	1,915	650	452	3,017
Graduate students	4,463	3,263	1,897	9,623
Undergraduate students	9,658	3,948	9,325	22,931
Total	17,676	8,695	12,233	38,604

C = control, T_1 = reminder and information on only two remaining clinics, T_2 = same information and call on self-restraint.
Tables.

Table 2. Randomization tests on treatment groups

	T_1 vs. C	T_2 vs. T_1	T_2 vs C
Professional categories	p-value for the test of equality between the randomized groups		
Faculty			
Gender	0.439	0.765	0.565
Age	0.163	0.163	0.343
Race	0.179	0.765	0.199
Occupation	0.378	0.042	0.132
Wage category	0.771	0.759	0.472
Staff			
Gender	0.403	0.724	0.220
Age	0.348	0.271	0.064
Race	0.832	0.679	0.704
Occupation	0.728	0.609	0.457
Wage category	0.303	0.775	0.692
Graduate students			
Gender	0.835	0.445	0.469
Undergraduate students			
Gender	0.593	0.268	0.137

The age distribution is given in 10 categories: less than 25 years old, 5-year intervals between 25 and 65 years old, and 65 years old and over. Races are white, asian, and others. Occupations for faculty are: tenured professors, non-tenured professors, recall and emeritii, lecturers, and others. Occupations for staff are executives and managers, professional staff, and support staff. Wage categories are: less than $40K, $40K-50K, $50K-60K, $60K-70K, and more than $70K.

Tests are Pearson's chi-squared corrected for the sampling design.
Source: Human Resources Customized Pivot Tables.

Table 3. Impact of information and self-restraint on demand: Average treatment effects

	Demand for flu vaccine			Impact of Information	Impact of Self-restraint	Impact of Info&restraint
	C	T_1	T_2	$T_1 - C$	$T_2 - T_1$	$T_2 - C$
	(st. error)	(st. error)	(st. error)	[t-stat]	[t-stat]	[t-stat]
Demand by professional category (in percentage of each category in the campus population)						
Faculty	3.0	6.5	3.8	3.4	-2.7	0.7
	(0.4)	(0.8)	(0.8)	[3.6]	[2.3]	[0.8]
% difference				112.4%	-42.0%	23.2%
Staff	1.3	1.8	2.4	0.6	0.6	1.2
	(0.2)	(0.5)	(0.7)	[1.0]	[0.7]	[1.6]
% difference				47.3%	31.8%	94.2%
Graduate students	0.7	1.2	0.8	0.5	-0.4	0.1
	(0.1)	(0.2)	(0.2)	[2.0]	[1.5]	[0.2]
% difference				61.6%	-33.8%	6.9%
Undergraduate students	0.4	1.1	0.5	0.7	-0.6	0.1
	(0.1)	(0.2)	(0.1)	[3.9]	[3.1]	[1.3]
% difference				169.1%	-50.9%	32.1%
Total demand (in % of campus population	0.8	1.6	1.0	0.8	-0.6	0.2
	(0.1)	(0.1)	(0.1)	[5.7]	[3.5]	[1.9]
% difference				110.1%	-37.5%	31.2%
Number of observations	147	149	98			

Source: Flu-shot survey, Fall 2004.

Table 4. Robustness test on average treatment effects

Dependent variable: Came to clinic A

	Constant	T_1	T_2	Male	Age category	Number of observations
Faculty	0.030	0.034	0.007	No	No	3033
	(0.005)	(0.012)	(0.012)			
	0.035	0.035	0.007	Yes	No	3033
	(0.010)	(0.012)	(0.012)			
	0.026	0.034	0.011	No	Yes	3033
	(0.035)	(0.012)	(0.014)			
Staff	0.013	0.006	0.012	No	No	3017
	(0.004)	(0.006)	(0.010)			
	0.015	0.006	0.011	Yes	No	3017
	(0.004)	(0.006)	(0.010)			
	0.016	0.007	0.012	No	Yes	3017
	(0.015)	(0.006)	(0.009)			
Graduate students	0.0074	0.0046	0.0005	No	No	9623
	(0.0015)	(0.0024)	(0.0028)			
	0.0086	0.0045	0.0007	Yes	No	9623
	(0.0018)	(0.0025)	(0.0026)			
Undergraduate students	0.0041	0.0070	0.0013	No	No	22931
	(0.0007)	(0.0031)	(0.0007)			
	0.0050	0.0071	0.0012	Yes	No	22931
	(0.0007)	(0.0031)	(0.0007)			

Standard error in parentheses, clustered at the department level.
Source: Flu-shot survey, Fall 2004.

Table 5. Heterogeneity in the impact of information and self-restraint on demand

	Number obs.	Demand for flu vaccine			Impact of		
		C	T_1	T_2	Information $T_1 - C$ [t-stat]	Self-restraint $T_2 - T_1$ [t-stat]	Info&restraint $T_2 - C$ [t-stat]
		(in percentage of campus population)			(in percentage difference)		
Member of a priority group	227	0.51	0.80	0.51	55.8	-35.8	0.0
					[2.4]	[2.9]	[0.0]
Previous users	190	0.43	0.66	0.42	54.0	-36.3	-1.9
					[2.2]	[2.7]	[0.1]
First-timers	31	0.06	0.13	0.09	120.6	-29.6	55.4
					[0.2]	[2.6]	[0.6]
Not member of a priority group	167	0.25	0.81	0.49	219.9	-39.2	94.3
					[3.6]	[3.4]	[2.3]
Previous users	80	0.10	0.41	0.26	318.8	-37.1	163.3
					[2.6]	[2.2]	[2.0]
First-timers	80	0.15	0.35	0.23	139.5	-34.4	57.2
					[2.1]	[1.8]	[1.3]

Source: Flu-shot survey, Fall 2004.

Table 6. Increase in demand induced by the shortage: Share of first-timers in 2002-2004 among participants

	Clinic A Campus group *C*	t-stat for difference with previous year[1]	Clinic A Non-campus community	t-stat for difference with previous year	Clinic B Non-campus community	t-stat for difference with previous year
2002	5.9		3.7		3.0	
2003	13.5	1.8	3.1	0.4	5.8	1.1
2004	25.7	2.9	11.9	4.0	22.6	4.7

[1]Standard errors clustered at the department*campus category level.

Source: Flu-shot survey, Fall 2004.

Table 7. Contrasting first-timers and previous users

	First-timers	Previous users	t-stat on difference
	(percent)		
Membership to official priority groups			
Adults 65 years of age or older	6.3	27.9	-2.8
Under chronic medical conditions	13.3	27.7	-1.8
Women who will be pregnant during the flu season	5.2	6.8	-0.3
Contacts with infant	6.8	5.5	0.3
Health-care worker	1.8	1.6	0.1
At least one of the above	33.4	64.2	-2.7
Other reason	68.7	44.5	2.4
Number of observations	112	466	
Reasons for wanting a flu shot among non members of priority groups			
Contact with children	4.0	17.9	-0.9
Can't afford to miss work or study	73.7	61.2	0.6
Believe shortage is just temporary	2.3	1.8	0.2
Concerned by shortage or potential epidemic	45.7	15.3	3.1
Other reasons[1]	17.4	43.3	-1.3
At least one of the above	91.4	88.1	0.3
Number of observations	57	69	

Non-campus community and campus groups C and T_1 (using sampling weight) from Clinic A.

Standard errors clustered at the department*campus category level.

[1] Other reasons include: living in dorms, being in contact with people, don't want to be sick, travel abroad.

Source: Flu-shot survey, Fall 2004.

Table 8. Evidence on effective screening, legitimate service, exclusion errors, and inclusion errors (cheating)

	Effective screening: Non-priority not serviced	Legitimate service: Priority serviced	Exclusion error: Priority not serviced	Inclusion error: Non-priority serviced	p-value for test of equality with group above	Share of cheaters among non-priority
	(1)	(2)	(3)	(4)	(5)	(6)
Criteria for definition of types						
Self-declared priority group	No	Yes	Yes	No		
Received flu vaccine	No	Yes	No	Yes		
	(Percent of participants in each category)					(%)
Clinic A: categories of participants						
Non-campus community	2.9	92.0	0.00	5.1		64.0
Campus group C	25.1	66.9	0.00	8.0	0.000	24.1
Campus group T_1	36.1	48.7	0.95	14.3	0.116	28.3
Campus group T_2	32.5	50.4	0.63	16.6	0.917	33.8
Clinic B: categories of participants						
Non-campus community	1.9	97.2	0.00	0.9		33.3
Campus population	6.8	88.2	0.79	4.2	0.002	38.1
Clinic A, campus group C						
Previous users	13.5	81.4	0.00	5.1		27.2
First-timers	56.7	28.8	0.00	14.5	0.000	20.4
Clinic A, campus groups T_1 and T_2						
Previous users	25.9	61.5	0.36	12.3		32.3
First-timers	50.7	25.8	1.92	21.7	0.003	30.0

Source: Flu-shot survey, Fall 2004.

Table 9. Impact of information and self-restraint on the number of vaccines distributed

	Number obs.	Number of vaccines distributed C	T_1	T_2	Impact of: Information $T_1 - C$ [t-stat]	Self-restraint $T_2 - T_1$ [t-stat]	Info&restraint $T_2 - C$ [t-stat]
		(in percentage of campus population)			(in percentage difference)		
Members of priority groups	224	0.51	0.78	0.51	52.8	-35.4	-1.3
					[2.3]	[2.8]	[0.1]
Non-members of priority groups	47	0.06	0.23	0.17	276.5	-27.6	172.6
					[1.9]	[1.1]	[1.6]
Difference					[1.5]	[0.3]	[1.6]
Total number serviced	271	0.57	1.02	0.67	76.6	-33.6	17.2
					[3.1]	[3.0]	[0.9]
Share of non-members of priority groups		0.11	0.23	0.25			

Source: Flu-shot survey, Fall 2004.

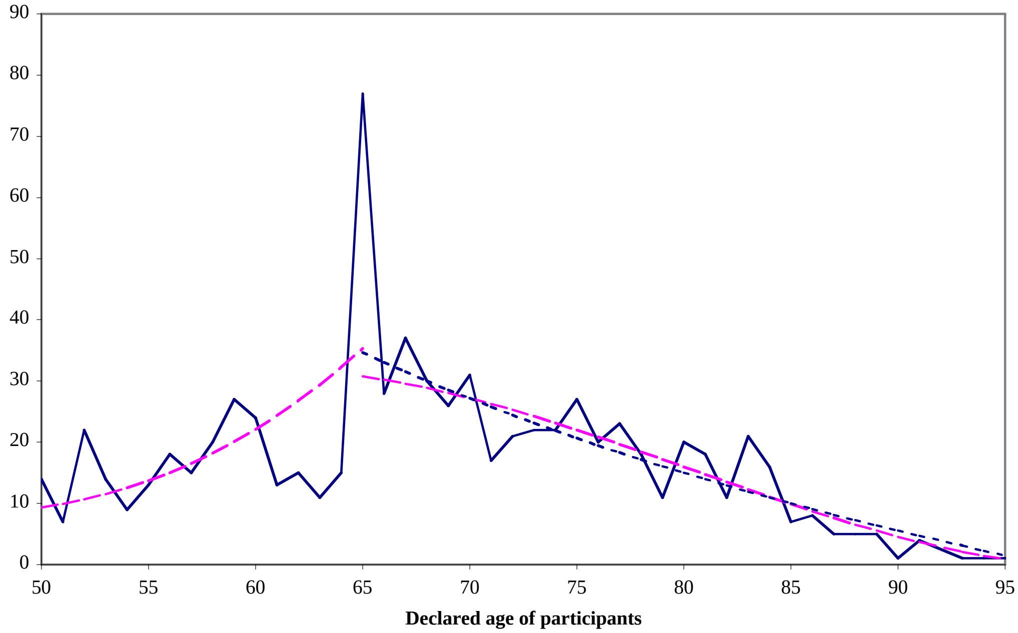

Figure 1. Evidence of cheating in self reporting: Bunching at age 65