

Ambiguity and Social Learning during the SARS Pandemic of 2003

Daniel Bennett, Chun-Fang Chiang, and Anup Malani

ABSTRACT. Epidemics such as swine flu generate a great deal of ambiguity and panic when they first emerge because there is little objective data on these risks. In this paper we examine the extent to which people turn to their peers, i.e., to social learning, for information on such new public health threats. Using microdata on the health care utilization during the SARS epidemic in Taiwan during 2003, we document that individuals protected themselves against contagion by avoiding medical facilities. We identify the effect of social learning on this avoidance behavior using predictions from a model of demand for medical care during an epidemic. The model implies that peers' beliefs are correlated with the change in peers' demand for medical care. Therefore we regress a patient's outpatient visits on the *change* in her peers' demand for medical care as a proxy for peers' beliefs. To rule out common unobservables, we control for the level of peers' demand and publicly reported SARS cases. We find that the elasticity of medical demand with respect to the beliefs of peers rose by 0.05 during SARS. Moreover, the role of social learning expended when actual cases declined at the tail end of the epidemic, suggesting that social learning may substitute for objective data on epidemic risk.

New epidemics have become common place. The last ten years have witnesses the global spread of Severe Acute Respiratory Syndrome (SARS), highly pathogenic avian influenza (H5N1), and swine flu (H1N1). Less-well reported are new local outbreaks of meningococcal disease and measles in Africa and foot and mouth disease in Europe. The toll from these epidemics depends in part on individual risk-taking behavior (Kremer 1996, Philipson 2000). That, in turn, depends on how people assess the threat from new epidemics.

Unfortunately new epidemics are difficult to evaluate. Often infections are rare, though outcomes can be severe. When HIV initially spread in the early 1980s, prevalence in the US was just a few thousand cases, while the life expectancy with AIDS was just 1-2 years (Osmond 2003).¹ Occasionally governments will distribute information on prevalence, but this data may not be timely or representative. For example, the US did not begin publishing HIV data until after 1985. Even then, it did not indicate whom it sampled or how many sexual interactions sample members had. As a result, individuals typically make decisions about risk-taking activities amidst a great deal of ambiguity about the probability and consequences of infection.

In this paper we explore whether individuals try to reduce this ambiguity through social learning. Specifically, do individuals augment information about new epidemics from public sources and individual observation with information acquired from neighbors? Our application is the 2003 SARS epidemic in Taiwan. Using panel data on the daily health care decisions of 5% of the population of that island nation, we study how individuals adjusted their risk taking in response to public reports of local SARS cases versus their neighbors' assessment of the risk of catching SARS.

We have three reasons for studying ambiguity and social learning in the context of epidemics. First, there is a growing interest in the role of social learning to explain a wide range of behaviors from technology adoption (e.g., Conley & Udry 2008, Bandiera & Rasul 2006, Foster & Rosenzweig 1995) to movie-going (Moretti 2009). A number of papers have specifically addressed health care consumption (Oster & Thornton 2009, Miguel & Kremer 2007, Deri 2005). It is natural therefore to consider whether social learning extends to individuals' assessment of new and significant public health risks. Our inquiry is most closely related to Rao et al. (2007), which examines how peers influence individuals' valuations of seasonal flu vaccine. The main differences are that Rao et al. study a small group of undergraduates and a recurrent disease. We shall examine the behavior of a large and diverse population in the context of an entirely novel epidemic.

Second, there is little formal connection between the literatures on ambiguity and on social learning. Much of the literature on ambiguity focuses on documenting ambiguity aversion (e.g., Ellsberg 1961, Einhorn & Hogarth 1986, Kunreuther et al. 1995) or theoretical modeling decision-making under uncertainty (e.g., Maccheroni et al. 2006, Klibanoff et al. 2005). There is some empirical work on how people try to reduce ambiguity. For example, Zimelman & Waller 1999 describe an experiment that shows individuals seek out more information – in their case they sample data on asset values – when

¹ Prevalence reached roughly 1 million in 1986, still just 0.3 percent of the total population. Outside of men having sex with men (MSM), among whom prevalence was 7.2% (Valleroy et al., JAMA, 2000), prevalence was substantially lower.

they confront ambiguity.² However, few papers have explicitly discussed the role of learning from peers in reducing ambiguity or of mimicry as a response to ambiguity. This connection is implicit, however, in the social learning literature that examines contexts steeped in ambiguity, such as new technology adoption.

We model risk-taking behavior using the smooth ambiguity decision model of developed in Klabinoff, Marinacci & Mukerji (2005) and show that learning – social or otherwise – can be driven by ambiguity aversion. The smooth ambiguity decision model represents preference by a nonlinear function over expected utility. The curvature of that function induces preferences over the variance of beliefs, which is how ambiguity is defined by Epstein-Zhang (2001) and Ghirardato-Marinacci (2002). A central result is that ambiguity increases demand for information to reduce variance of estimates about the probability of different health states. This highlights an important difference between our study and prior papers on social learning in ambiguous contexts. Social learning in the context of technology adoption might serve to both correct negative bias in assessments of technology as well as increase precision of those estimates. The bias correction is evidenced by the gradual (rather than sudden) adoption of technology. In our application, individuals likely had access to an unbiased, though perhaps imprecise, estimate of SARS prevalence. To the extent that they used social learning to improve their information set, their goal was more plausibly to improve precision due to ambiguity aversion.

Third, social learning has serious, negative consequences for the spread of disease. Whether and how fast a disease spreads within a population depends the number of individuals that a sick person infects. If the number is below one, the disease will die out. However, if there is heterogeneity in this number across the population, the disease may survive among high risk individuals. A disease will be even more resilient if high risk individuals are geographically clustered. This may occur if high risk individuals seek each other out, as might be the case with sexually transmitted disease (STDs) such as HIV (Kremer 1996). But the phenomenon extends to all communicable diseases if individuals who happen to be neighbors develop correlated beliefs about the disease and thus risk behaviors. Social learning may be the mechanism generating these correlated beliefs.

With this motivation we turn to our application. SARS spread to Taiwan from mainland China in March 2003. By July, when the outbreak was contained, it had infected 671 individuals and killed 84. Although at its height SARS infected 60 individuals in one day and the case fatality rate was 9.6%, overall prevalence was only 3.7 per 100,000 population. The public sought to protect themselves against SARS by avoiding public spaces, especially hospitals and medical clinics. This illustrated in Figure 1A, which documents weekly SARS cases, and Figure 1B, which plots the decline in medical visits in 2003. The response was stark: visits fell by 20% during the epidemic.

² Eliaz and Schotter (2007) also describes an experiment where subjects pay for information that does not change their choices but does improve confidence about their choice. They interpret this result as evidence of a preference for confidence, and reject that it reflect a preference for early resolution of uncertainty (Kreps & Porteus 1978). But a preference for confidence may also be described as a static preference for more certainty when probabilities are subjective.

Of particular interest is the sustained reduction in visits, even after SARS cases disappeared. Surely individuals were concerned about SARS reappearing. During this period, when the publicly reported risk was constant – at zero – across townships, there remained, however, a large degree of cross-sectional variation in visit rates. This is illustrated in Figure 1C, which plots the coefficient of variation³ in visits in 2002 and 2003. The coefficient of variation jumps during the SARS period in 2003 to 4.6. By comparison the coefficient of variation is just 3.9 during the same period in 2002. This excess variation in visits suggests a role for social interactions in decision-making about medical visits.

The central challenge with identifying the influence of peer effects on behavior is to distinguish correlated behavior induced by common but unobservable factors from that induced by the sharing of information. This is the so-called reflection problem (Manski 1993). The prior literature on peer has taken a number of different approaches to address this problem, from randomizing information to different social networks (e.g., Duflo et al. 2009, Sacerdote 2001) to identifying separate spaces through which common unobservables and information spread (e.g., Conley & Udry 2009, Munshi & Myaux 2006). Like Moretti (2009), our approach is to use restrictions from an economic model – of the demand for medical visits – to distinguish peer effects from other common unobservables.

Our model assumes that the supply of visits is infinitely elastic so that the equilibrium quantity of visits is determined by the intersection of the government-set price of visits (co-pay) and the demand curve for visits. SARS acted like a tax on visits, reducing equilibrium quantity. The magnitude of the tax depended on individuals' assessments of the mortality risk from SARS and the value of life. These assessments could be backed out with information on the change in visits after SARS and the slope of the demand curve. If one assumes that peers have similar demand curves and value of life, then one can identify whether one individual's assessment of risk depended on their neighbor's assessment if the former's visit behavior depends on *changes* in the latter's visit behavior.

Therefore, we regress an individual's weekly visits to a doctor on the publicly reported level of SARS cases, the level of visits by peers to the doctor, and the year-on change in the level of visits by peers to hospitals. The number of publicly reported SARS cases provides an unbiased estimate of the current prevalence of SARS in the population. The level of peer visits captures the number of contacts an individual can expect with other people who may be carrying SARS. It also captures common unobservables, such as supply shocks.⁴ This leaves the change in the level of visits to capture the effect

³ We report coefficient of variation rather than simple variance because SARS reduced visits and visits are bounded below at zero. Average weekly visits before SARS were close enough to zero that the reduction in visits due to SARS mechanically reduced the variance of visits. The coefficient of variation addresses this by normalizing the standard deviation of visits by the mean of visits. It is unlikely that the elevation in the coefficient of variation is caused simply by normalization with the mean. The reason is that the coefficient of variation is larger during SARS than the Chinese New Year, yet mean visits were lower during the Chinese New Year than during SARS. Everyone stopped going to doctors during the New Year holidays, but there was variation in the reduction in visits during SARS.

⁴ The level of peer visits may also capture some peer effects. We do not rely upon it for identification of peer effects because it also captures common unobservables. Instead we rely upon the change in peers' visits conditional on the level of peers' visits to capture peer effects. Since the change in peers' visits may not capture all peer effects, it provides a conservative test of social learning.

of peers' assessments about SARS risk. Identification is strengthened by the use of cross-sectional fixed effects to capture time-invariant unobservable effects and an interaction with a fixed effect for the entire SARS period to capture unobservable changes during the SARS epidemic.

We draw two conclusions from our analysis. First, peer's assessments of SARS risks had a significant influence on individuals' visit behavior. Their effect, although small in absolute value, implies an elasticity of individual medical demand with respect to peer's beliefs of 0.05.⁵ This is about half the elasticity with respect to other measures of SARS risk, such as newspaper reports about the number of local SARS cases and information about how crowded waiting rooms are, and other common unobservables correlated with actual SARS risk, such as closures of doctor's offices. Thus social learning is about half as influential as other, largely non-peer sources of information. Second, peers had a greater effect when individuals lacked positive reports of local SARS cases. We find that peer effects are smaller during the height of the SARS epidemic than in the period just after it, when patients had no formal reports of SARS cases but might have believed the epidemic risk had not completely dissipated. Likewise, peer effects were smaller when a patient's township was currently experiencing a SARS cases than after her township had already experienced a case.

The remainder of this paper may be outlined as follows. Section I describes the SARS epidemic in Taiwan and the precautions residents took to avoid SARS. Section II presents a model of the demand for medical visits during the epidemic. Section III elaborates on our empirical strategy for separating peer effects from common unobservables and describes our data. Section IV presents our empirical results.

I. SARS IN TAIWAN

SARS is an acute respiratory infection, caused by a coronavirus, that initially resembles a cold but then progresses into severe pneumonia. The SARS epidemic, which took place from November 2002 to July of 2003, sickened 8096 people and killed 774 globally. The outbreak began in the Guandong, China, which is the province adjacent to Hong Kong. Epidemiologists trace the disease to human interaction with the civet cat, an exotic species that residents of Guandong cultivate as a delicacy. After achieving human to human transmission, the SARS epidemic quickly spread to Hong Kong and then to other parts of Southeast Asia. Rigorous quarantine and infection control measures by the WHO are credited with containing the epidemic during the summer of 2003.

SARS reached Taiwan relatively late in the progression of the epidemic. The first cases arrived in March of 2003 among travelers who had visited mainland China. While these cases raised concern about an indigenous outbreak in Taiwan, they did not cause excessive alarm. On April 24, 2003, a cluster of infections at the Ho-Ping Hospital in Taipei began a chain of local transmission that led to 116 probable cases and 10 deaths in the ensuing two weeks. At the epidemic's peak in late-May there were

⁵ More precisely, we estimate that a 1% increase in our proxy for peer's beliefs about SARS risk increases visits over several periods. The sum of these increases is equivalent to a 0.5% change in one period of visits.

over 60 reported SARS cases in a single day. By the end of the epidemic in early July, 671 confirmed cases had led to 84 deaths on the island (Hsieh, Chen & Hsu 2004).

The SARS epidemic caused panic in Taiwan. Surveys reported intense anxiety and insecurity about the SARS among the population (Ko et al. 2006). People avoided public places and wore facemasks whenever they ventured outdoors (Institute of Sociology-Academia Sinica 2003).⁶ Commerce slowed as people avoided transactions that could potentially expose them to SARS, such as restaurant meals.

SARS acutely impacted the health care system. Since most SARS cases resulted from transmission in the health care setting (Hsieh, Chen, & Hsu 2004), many people perceived an increased risk of SARS from visiting hospitals and clinics (Institute of Sociology, Academia Sinica 2003). A wholesale decline in health care utilization began at the time of the Ho-Ping outbreak and lasted through September, long after the final SARS cases. During this time, outpatient visits declined by 25 percent nationwide (Lu, Chou & Liou 2007).

Panels A and B of Figure 1 illustrate the time path of the SARS outbreak and the related decline in health care attendance. These figures are labeled according to three distinct SARS periods, which are described below. Figure 1A plots the number of reported and probable SARS cases by week during 2003. The classification criteria for reported cases are less stringent than for probable cases. The figure shows the period of sporadic infection during March and early April, followed by a period of peak infection, and then a subsequent decline. Figure 1B plots the total number of outpatient visits in Taiwan by week for both 2002 and 2003. The 2003 series shows the precipitous decline in visits that coincides with peak SARS period. Visits decline by 36 percent from Week 11 to Week 23. The level of visits does not fully recover until September. The series for 2002 provides a point of comparison. While both years exhibit seasonal fluctuations and a decline that coincides with Chinese New Year, there is no precipitous drop in visits during the late spring of 2002.

The magnitude and duration of this trough are puzzles in light of existing research about the prevalence response to disease risk. While SARS had a high case fatality rate (9.6%), prevalence was low even during the peak period (3.7/100,000 population). That such a small risk could cause such an outsized response suggests a high degree of risk aversion or perhaps ambiguity aversion, particularly in areas where SARS was never reported.⁷ The extended duration of the trough also suggests that people

⁶ The Institute's Special Survey on SARS from May 2003 reported that, when asked what they did to prevent SARS, 34% of people responded that they avoided public places.

⁷ Quarantines of infected patients and individuals in close contact with infected patients were used to contain SARS (Hsieh et al. 2005, CDC Taiwan 2003). Although a total of 151,460 individuals were confined, these quarantines cannot explain the decline in visits. First, individuals were only quarantined for 10-14 days. This period is shorter than the 3 week median time between medical visits among the population. Second, quarantines required patients to go to the hospital for any respiratory or cold symptoms. Third, 95,828 of those quarantined were under Level B quarantine that permitted individuals much freedom of movement, including visits to the doctor. Moreover, this category mainly included foreign travelers, many of whom would not have been generating demand for visits in 2002, the control year. Third, the prevalence of confirmed SARS cases among those quarantined (11/100,000) is much higher than among those in the population (0.37/100,000). Thus SARS cases

may have lacked information on SARS risk even after it vanished. This paper investigates the hypothesis that social learning contributed to this SARS response. Lacking informative data about the true risk of SARS, people observed their peers to formulate perceptions of disease risk.

II. MODEL OF THE DEMAND FOR INFORMATION AND DOCTORS' VISITS

To motivate our strategy for identifying peer effects and to connect our application to the literature on ambiguity, this section develops a model of individual demand for information and medical visits in the presence of SARS. We will begin with a description of expected utility from sickness and from medical care. Medical care eliminates sickness but, during the SARS epidemic, involves a risk of catching SARS. We permit individuals to gather information on the risk of SARS, but at a cost. This information may be used to reduce the variance of patients' estimates of SARS risk. Finally we will model how demand for information and visits changes when we introduce ambiguity aversion using the smooth ambiguity preferences of Klibanoff et al. (2005).

Our analysis yields two conclusions. First, SARS functions like a tax on the price of medical care that reduces the equilibrium levels of medical visits. We will use this fact to justify our strategy for identifying peer effects in Section III. Second, ambiguity aversion generates demand for improvements in the precision of patients' estimates of SARS risk. While there may be demand for information about SARS in the absence of ambiguity, it is amplified by the presence of ambiguity. To the extent that the SARS epidemic created ambiguity, one might infer that additional demand for information about SARS was partly driven by that ambiguity. Indeed, if patients were only searching for information about the probability of catching SARS and received unbiased news reports about the risk of SARS, one could conclude that ambiguity alone explained the additional search for information – including from peers – during the epidemic.

A. Demand for medical care

A patient obtains utility from wealth and health. For simplicity we assume consumption of these two goods are additive. Thus, a patient's utility when sick is $u(w, h - s) = u(w + h - s)$ where w is wealth, h is baseline health, and $s < h$ is the level of sickness. We assume that $u(0) = 0$, $u' > 0$ and $u'' < 0$. If we assume that each individual may visit a doctor only once per period and that medical care perfectly cures sickness (other than SARS), then the patient's utility from medical care is $u(w + h - c)$ where c is the co-pay for a medical visit. Let q be the probability of catching SARS at a medical facility and dying. For simplicity, we ignore the morbidity from SARS. If we assume that utility $u(w, 0) = 0$ when the patient is dead, the patient's expected utility from visiting the doctor during the SARS epidemic is $qu(w - c, 0) + (1 - q)u(w - c, h) = (1 - q)u(w + h - c)$.

The government of Taiwan pays the cost of all health care above the co-pay. This payment, which resembles cost-plus reimbursement, is sufficiently generous to encourage hospitals to bid for

and quarantines are positively correlated. Our regression indirectly controls for quarantines by controlling for SARS cases.

patients (Cheng 2003). Therefore, we assume the supply of medical care is perfectly elastic at the co-pay. This implies that shifts in the equilibrium level of visits can identify shifts in the demand curve. It is possible that, in reality, supply was not perfectly elastic due to SARS-related supply shocks. For example, doctors may have closed their offices. In the empirical analysis we attempt to control for supply shocks or assume that the patients observe supply shocks as well as the econometrician so they are able to back out shifts in the demand curve from changes in equilibrium quantity.

SARS suppresses the equilibrium level of medical visits in much the same way a tax on copay would. Before the SARS epidemic, a patient would visit the doctor if the co-pay was less than the cost of sickness: $u(w + h - c) > u(w + h - s)$ or $c < s$. This is because we have assumed that medical care is perfectly curative. If patients differed in the severity of their sickness, the aggregate demand for medical visits would be $D(c) = N \int_{s>c} f(s)ds$, where $f(s)$ is the distribution of sickness in the population and N is the population of patients. This is also the equilibrium level of visits since supply is perfectly elastic. During the SARS epidemic, however, a patient would visit the doctor only if $(1 - q)u(w + h - c) > u(w + h - s)$. Because SARS lowers the utility of visiting the doctor by $(1 - q) < 1$, an individual has to be more sick during SARS before she would go to the doctor. This will lower demand for medical care.

This shift in demand is equivalent to holding demand constant but imposing a tax on co-pay. Since utility comparisons are preserved under monotonic transformations, the visit condition during SARS is equivalent to $\ln(1 - q) + \ln u(w + h - c) > \ln u(w + h - s)$. If we take a first-order Taylor approximation of $\ln u(w + h - x)$ around $(w + h)$, this condition can also be written, for small s and c , as

$$\ln u(w + h) - R(w + h)s < \ln(1 - q) + \ln u(w + h) - R(w + h)c$$

where $R(w + h) = u'(w + h)/u(w + h) > 0$.⁸ Solving for s , the visit condition becomes $s > c + T$ where

$$T = -\ln(1 - q) / R(w + h) > 0 \tag{1}$$

Thus, aggregate demand for medical visits falls to $D(c + T) = N \int_{s>c+T} f(s)ds$. Since supply is perfectly elastic, this effect of this shift in demand on equilibrium quantity is equivalent the effect of imposing a tax of T on copay but leaving demand at the pre-SARS level. This is illustrated in Figure 2. We can see from (1) that this tax is proportional to the probability q of catching SARS.

B. Demand for information on SARS

We now introduce information-gathering into the model. We assume that individuals do not know the true probability q of catching SARS. Instead they obtain a costless public signal $\hat{q} = q + e$ from local media, where e is mean-zero error in the public signal and is orthogonal to q . In addition,

⁸ Because utility is unique to affine transformations, this result does depend on utility $u(x)$ being positive. Negative utility can always be transformed into positive utility.

individuals can purchase n additional “data points” on SARS risk at marginal cost p . These data points could be time spent researching SARS online and at the library. It could also be conversations with neighbors and friends discussing the risk of SARS, so long as individuals believe these are not merely reflections of the public signal, but actually contribute new information about SARS. The additional information may be used to reduce the variance of the signal. We capture these benefits by assuming that e has variance $\sigma^2/(n + 1)$. Thus an individual’s posterior $\mu(q)$ on the probability of catching SARS has the following moments

$$E_{\mu}[q] = E + E_{\mu}[\hat{q} - e] = \hat{q}, \quad Var_{\mu}[e] = Var_{\mu}[\hat{q} - e] = \sigma^2/(n + 1)$$

We define uncertainty as variance in posteriors about a variable. This is consistent with the approach of Epstein-Zhang (2001) and Ghirardato-Marinacci (2002). In our model, the only source of uncertainty – variance in posteriors – is the probability of catching SARS. There is no uncertainty about arguments in the utility function, such as wealth or the co-pay.

In a subjective expected utility framework, risk aversion only generates demand for information to reduce uncertainty when there is uncertainty about consumption in states, not when there is uncertainty about states. Uncertainty about the probability of states generates no demand for information to reduce that uncertainty because expected utility is linear in probabilities over states. If instead there were uncertainty about arguments in the utility function, subjective expected utility would generate demand to reduce uncertainty because the utility function is concave.⁹

Therefore, subjective expected utility does not generate demand for information about SARS. This is illustrated by a patient’s subjective expected utility from visiting a medical facility during the SARS epidemic,

$$E[1 - q|\hat{q}, n]u(w + h - c - pn) = (1 - \hat{q})u(w + h - c - pn)$$

Purchase of additional data points has marginal cost p but offers no marginal benefit to subjective expected utility. It is true that additional data points lower the variance $\sigma^2/(n + 1)$ of a patient’s posterior about SARS risk, but subjective expected utility from visiting the doctor depends only on the mean of the patient’s posterior, not its variance.

Our claim that mere risk aversion cannot explain demand for information during SARS rests on four important assumptions. The first is our use of the subjective expected utility framework. If, instead, one employed an objective expected utility framework, then even uncertainty about consumption in states might not generate demand for information because uncertainty might not imply variance in – and thus risk from – objective variables. For example, uncertainty in information about wealth or co-pay does need not imply variance in actual wealth or co-pay. Our use of subjective expected utility to model risk aversion should not be controversial. It is certainly more widely accepted in economics than objective expected utility. Moreover, any decision-making framework that introduces curvature in preferences over probabilities of states not only has risk aversion, but also

⁹ In theory, uncertainty about the shape of the utility function itself could also generate demand for information to reduce that uncertainty.

ambiguity aversion. We shall illustrate this in the next section, where we argue that ambiguity aversion can generate demand for information even when there is uncertainty only about probabilities of states.

Second, we assume that uncertainty during SARS affected only the probability of different states of the world. This is not an unreasonable assumption. Even if there were uncertainty about the effects of SARS, e.g., whether it caused death, hospitalization, or merely some sick days, this could be expressed in the model by adding to the two states we have – no SARS or SARS and death – additional states such as SARS plus hospitalization. We cannot think of an important source of uncertainty introduced by the SARS epidemic that requires expression in the utility function.¹⁰

Third, we assume that the public signal provides an unbiased estimate of SARS risk. This too is a reasonable assumption. Having compared reported SARS cases and confirmed SARS cases, we see no bias in news paper stories about reported cases. Nor do surveys about media performance during the epidemic suggest impressions of bias among respondents (Lu 2006). Even if there were bias, this would not be problematic for our claim so long individual's had defined, finite expectations about that size and sign of bias. In that case, individuals would simply subtract expected bias from the mean of the public signal to get an unbiased estimate.¹¹

Fourth, we assume that individuals do not have preferences over error in the public signal about the probability of catching SARS. This is how other papers using an expected utility framework have generated demand for information. For example, Moretti (2009) generates social learning by assuming people hate to be surprised about movies. Error might also matter if there are adjustment costs to correcting a bad decision. This may explain the demand for social learning in papers about technology adoption. Finally, only might imagine regret generates preferences over error. We do not think these or other possible consequences of error have first order effects on patient behavior during SARS.

C. Decision making under ambiguity

Although uncertainty about the probability of SARS does not generate demand for information when patients are merely risk averse, it can generate demand for information if patients are ambiguity averse. We model this possibility using the smooth ambiguity preferences (SAP) of Klibanoff et al. (2005). The SAP framework adds two features to expected utility. First, ambiguity is defined as uncertainty about the probability distribution function over states. In our model of SARS risk, this uncertainty is captured by the posterior distribution $\mu(q)$ over the different probabilities q of catching SARS and dying. Second, in the presence of ambiguity, individuals do not choose actions directly on the basis of expected utility. Instead, they act on the basis of expectations about a smooth function ϕ over expected utility. If this function is nonlinear, an individual's preferences are no longer linear in the

¹⁰ An example, not associated with the epidemic, would be uncertainty about the current value of one's investment portfolio. Measurement error in the value of that portfolio does not represent actual risk in the portfolio, but rather uncertainty about valuation of that portfolio.

¹¹ If individuals did not have finite expectations about bias in the public signal, this would pose problems not so much for our claim about demand for information under subjective expected utility, but about subjective expected utility itself. With undefined expectations about bias, individuals could not calculate and thus act upon subjective *expected* utility.

probabilities of different states; rather the individual cares about other moments – importantly, the variance – of those probabilities. If ϕ is concave, then an individual is averse to variance in possible distributions of states, and thus to ambiguity.

For convenience we assume $\phi(x) = \ln(x)$ so that the overall utility from visiting the doctor during the SARS epidemic is

$$\int \ln(1 - q)u(w + h - c - pn) \mu(q) dq = E_{\mu}[\ln(1 - q) | \hat{q}, n] + \ln u(w + h - c - pn) \quad (2)$$

Further, we assume e is log-normal,¹² so that the overall utility can be written as a simple function of the moments of e :

$$E_{\mu}[\ln(1 - q) | \hat{q}, n] = \ln(1 - \hat{q}) - \frac{1}{2} \left\{ 1 + \frac{\sigma^2/(n + 1)}{(1 - \hat{q})^2} \right\} \quad (3)$$

Our assumption about functional forms of ϕ and e are not necessary to demonstrate that variance $\sigma^2/(n + 1)$ in information about SARS risk reduces utility, but they do help clarify the relationship.

In our simplistic model, the patient does not risk catching SARS if she does not visit the doctor. (This is unrealistic, but will not affect our basic point that ambiguity aversion generates demand for information.¹³) Therefore the patient has no value for information on SARS in the case where she chooses not to visit the doctor. This inference, when combined with the fact that, in our model, gathering data has a determinative effect on the variance of posteriors about the public signal, suggests a patient can choose how much information to gather and whether to visit the doctor in a sequential manner. First, focusing on the case where she visits the doctor, the patient determines optimal n^* . Then she chooses between visiting the doctor (and gathering n^* data points) and staying home (and gathering no data).

The patient chooses how much information to sample (beyond the public signal) by maximizing the overall utility from visiting the doctor (2) over n . The first order condition is

$$\frac{1}{2} \left\{ \frac{\sigma^2/(n^* + 1)^2}{(1 - \hat{q})^2} \right\} = p \frac{u'(w + h - c - pn^*)}{u(w + h - c - pn^*)}$$

The left hand side is the marginal benefit of reducing the variance of beliefs about the risk of SARS. The right-hand side is the marginal cost of information – the price of information times the marginal utility of consumption. The marginal benefit of costly information increases in the variance of the public signal about SARS. Thus the demand for information increases in the level of ambiguity about SARS risk:

¹² Obviously a log-normal variable cannot have mean zero. We can address this by assuming that e has mean m and that the public signal $\hat{q} = q - m + e$. This adds extra terms to the analysis, but does not affect our results. So we proceed as if e remains mean-zero.

¹³ It is unrealistic on its face and by its implication that people do not seek out more information if they do not go to the doctor. This is not critical because a model where individuals chose n before or at the same time as they chose whether to visit the doctor would also generate a demand curve that increased in the variance of the public signal.

$$\frac{dn}{d\sigma^2} = \frac{-(1/H)}{(n+1)^2(1-\hat{q})^2} > 0$$

where $H < 0$ is the second derivative of overall utility with respect to n . So long as $\sigma^2/2(1-\hat{q})^2 > pu'(w+h-c)/u(w+h-c)$, the marginal benefit of gathering information is greater than its marginal cost at $n=0$, so the patient will demand more information under smooth ambiguity preferences with ambiguity aversion than under subjective expected utility and mere risk aversion.

It appears from our model that the individual demand for SARS does not depend on how sick and individual is. This is not true if individuals care about SARS risk even when they do not visit a medical facility because sickness affects the marginal utility of consumption or if medical care does not perfectly cure sickness. But if we proceed with our simple model, demand for information to supplement the public signal is simply n^* times the demand for visits.

The demand for visits, however, is lower when patients are averse to ambiguity. We can see this by converting the demand for visits into an equivalent tax on co-pay and comparing the tax under ambiguity aversion to the tax T under simple risk aversion reported in (1). Under ambiguity aversion a patient visits the doctor if

$$\ln(1-\hat{q}) + \frac{1}{2} \left\{ 1 + \frac{\sigma^2/(n^*+1)}{(1-\hat{q})^2} \right\} + \ln u(w+h-c-pn^*) > \ln u(w+h-s)$$

The left hand side is derived by plugging (3) into (2). Taking a first order approximation of $\ln u(w+h-x)$ around $(w+h)$ and solving for s , we find that a patient visits a medical facility only if $s > c + T + A(\sigma^2)$ where

$$A(\sigma^2) = \frac{1}{2R(w+h)} \left\{ 1 + \frac{\sigma^2/(n^*+1)}{(1-\hat{q})^2} \right\} + pn^* > 0$$

where, recall, $R(w+h) = u'(w+h)/u(w+h) > 0$ is the marginal utility of consumption at $w+h$. As with the individual demand for information, the additional tax from SARS under ambiguity aversion is increasing the variance of the public signal.¹⁴ The same is true for aggregate demand, $D(c+T+A(\sigma^2)) = \int_{s>c+T+A(\sigma^2)} f(s) ds$, and, because supply is perfectly elastic, the equilibrium level of visits.

Finally, increasing the probability of SARS – as opposed to uncertainty about SARS – reduces both subjective expected utility and over all utility under smooth ambiguity preferences. However, ambiguity aversion amplifies risk aversion so an increase in risk of SARS has a bigger impact under smooth ambiguity preferences than under subjective expected utility:

$$\frac{dE_{\mu}[\ln(V)]}{d\hat{q}} = \frac{dV}{d\hat{q}} + R \frac{dA}{d\hat{q}}$$

¹⁴ Sampling more data does not increase the implied tax from SARS. The reason is that A is decreasing in data points when $n < n^*$ because n^* maximizes overall utility by minimizing A .

Where V is subjective expected utility from visiting the doctor during SARS and $dA/d\hat{q} > 0$.

III. EMPIRICAL STRATEGY AND DATA

A. Empirical strategy

It is our contention that, when the SARS epidemic began, patients were confronted with substantial ambiguity. Although they received a daily public signal about the risk from SARS, they sought out more information about the probability of catching SARS. In particular, they sought information from their peers. Our empirical analysis seeks out evidence of this social learning.

We begin with an empirical model of visits of the form

$$v_{ijkt} = \alpha + \beta E[q_{ijkt}] + u_{ijkt} \quad (4)$$

where v_{ijkt} is the number of visits patient i in location j pays to doctor k during week t and $E[q_{ijkt}]$ is the patient's expectation about the probability of catching SARS at that doctor's office and time. This model can be derived from the demand for visits derived in Section II.¹⁵ We do not observe expectations about q_{ijkt} , so we must proxy for it. However, we are not interested in the effect of expectations on visits so much as the effect of peer behavior on expectations. Therefore, we shall estimate a reduced form of (4) where visits are regressed directly on proxies for expectations.

1. Identification of peer effects

Our data on the behavior of medical patients in Taiwan are not sufficiently fine grained to separate out individuals who are in the same social network but not subject to common unobservable shocks. Therefore, we cannot use the correlation between the behavior of individuals and the behavior of their social network to identify peer effects. Instead, we propose to use our model of visits to identify correlations in the behavior of peers that can only be explained by social learning.

The social learning we focus on is either observation of the medical visit behavior of peers or conversations with peers about the health risk from SARS. According to our model, the expected risk from SARS operates on the equilibrium level of visits like a tax on the copay for a visit. One can back out the expected risk from SARS with knowledge of the change in visits due to SARS and the slope of the demand curve for visits. Indeed, if the demand for medical visits, aside from the influence of SARS, is stable over time, the change in visits due to SARS will be proportional to the expected medical risk from SARS. Therefore, we proxy the effect of peer's beliefs about the SARS risk with the year-on change in the average number of visits by peers: $\Delta \bar{v}_{jkt} = \bar{v}_{ijkt} - \bar{v}_{ijkt-52}$, where $\bar{v}_{ijkt} = \sum_{i' \neq i} v_{i'jkt} / (N_{jkt} - 1)$ and N_{jkt} is the number of patients that regularly see doctor k .

¹⁵ The average level of visits is $D(c + T + A)/N$. If we take a linear approximation around c , average visits is $(1/N)D(c) + (1/N)D'(c)(T + A)$. Substituting for $T + A$ we may write this quantity as $\alpha_0 + \alpha_1 A(E[q]) - \beta \ln(1 - E[q])$. Since $E[q]$ is very small, $A \approx (1/2)R(w + h)\{1 + \sigma^2/(n^* + 1)\}$, which does not depend on q , and $\ln(1 - E[q]) \approx -E[q]$.

Intuitively, our method of proxying for social learning with changes in equilibrium demand for visits identifies peer effects under either of two conditions. First, the patient converses with peers about the risk from SARS but the econometrician does not observe this. If supply is perfectly elastic and there are no SARS-related supply shocks, then the peer's beliefs will be reflected by a change not just in their demand but also in equilibrium quantity. If the econometrician observes this change in equilibrium quantity, he possesses a variable that is proportional to the change in demand and thus peers' beliefs about SARS risk.¹⁶ Thus a change in equilibrium quantity is a valid proxy for peers' beliefs. If there are supply shocks, the change in quantity remains a valid proxy so long as the econometrician can control for supply shocks.

Second, the patient does not converse with peers, but does observe the change in peers' visit behavior over time. Medical facilities are located within the community and people have the opportunity to observe directly whether other people are utilizing these facilities. If the econometrician also observes this change in visit quantity, then he too observes the patient's social learning. If the patient can identify the change in demand from the change in quantity, e.g., because she has knowledge of the slope of peers' demand curves, then the econometrician has only a variable that is proportional to the patient's social learning.

2. Common unobservables

Having a proxy for peer beliefs is a necessary but insufficient condition for identifying peer effects. Peer beliefs may respond to a common but unobservable shock, such as the underlying risk of SARS or a common shock to the supply of medical visits. Therefore, our model includes a number of controls for such common unobservables.

First, we include other measures of the underlying probability of catching SARS during a visit to the doctor. This probability depends on the number of individuals a patient will come into contact with at a medical facility and the probability of catching SARS given a contact. Patients learn about the probability of catching SARS given a contact from the public signal on SARS: daily news reports of the number of people s_{jt} in location j infected with SARS at time t .

Patients learn about the number of individuals they will encounter in the waiting room by observing total number of visits to the doctor by one's peers, $\Sigma v_{jkt} = \Sigma_{i' \neq i} v_{i' jkt}$, or the average number of visits to the doctor by one's peers, $\bar{v}_{jkt} = \Sigma v_{jkt} / N_{jkt}$, and the number of other patients N_{jkt} that visits the same doctor's office. This can be observed directly during a prior visit, ascertained by walking by the medical facility, or by calling the doctor's office. While this term is correlated with our proxy for social learning, it is not the same as that proxy. The difference is that $\Delta \bar{v}_{jkt}$ includes the

¹⁶ The econometrician cannot identify peers beliefs without more information, including the slope of peers' demand curves. But a proxy requires something positively correlated with

previous year’s level of visits. Without that information, one cannot identify the *change* in visits and thus the tax implied in peers’ beliefs about SARS risk.¹⁷

The level of visits by peers not only captures the underlying rate at which a patient will contact other people at the doctor’s office, it also captures supply shocks that operate at the level of the medical facility. For example, if the doctor requests that patients reduce their level of visits, this will appear in the total number of visits, but not in the change in visits.¹⁸

Second, we include week fixed effects and both individual and doctor’s office¹⁹ fixed effects. The time fixed effects capture common beliefs – at the national level – about the severity of SARS or changes in the national government’s policy about reporting or battling SARS. The individual and doctor’s office fixed effects capture time-invariant common unobservables. This includes heterogeneity, for example, in long-term medical conditions that may lead to variation in attendance or in risk aversion that could lead to individuals sorting to different townships or doctors.²⁰ Since there is only cross-sectional variation in the number of patients N_{jkt} that visit each doctor’s office, this variable will be dropped from regressions in lieu of the doctor’s office fixed effects.

Finally, we want not merely to identify social learning, but rather to identify social learning induced by the SARS epidemic. Therefore, we interact our measure of social learning, $\Delta\bar{v}_{jkt}$, and our measures of the total number of visits, Σv_{jkt} , with an indicator S for the period during which Taiwan experienced the SARS epidemic. This structure also helps rule out the influence of additional common unobservables. Even if common unobservables create a correlation between visits and our measure of social learning in general, they must differentially cause this correlation during the SARS period specifically to bias our estimates of social learning during the SARS period.

3. Baseline specification

These considerations yield the following regression model:

$$v_{ijkt} = \beta_1 \bar{v}_{jkt} + \beta_2 \Delta\bar{v}_{jkt} + \beta_3 S_{jt} + \beta_4 S \bar{v}_{jkt} + \beta_5 S \Delta\bar{v}_{jkt} + \alpha_i + \delta_j + \gamma_t + \varepsilon_{ijkt} \quad (5)$$

We use \bar{v}_{jkt} rather than Σv_{jkt} in the baseline specification because \bar{v}_{jkt} better captures all common unobservables that might affect the average number of peer visits each week and thereby ensures that

¹⁷ We do not claim that total visits by peers does not pick up any peer effects. Our claim is that it may also pick up common unobservables like the number of contacts in the doctor’s office. In short, although it does not cleanly identify peer effects, it does control for common unobservables

¹⁸ Unless the doctor rations visits in a pro rata manner, e.g., by telling all patients to visit 20% as much as they visited last year, this doctor-level supply shock is unlikely to be captured by our proxy for peer effect, $\Delta\bar{v}_{jkt}$. However, we will address this problem by looking at the behavior of different peer groups within a given doctor’s patient pool.

¹⁹ As we shall discuss in section III.B, we define social networks at the doctor-facility level. Since doctors may work at different facilities, and have different populations of patients at each facility, we have separate fixed effects for each doctor-facility pair, i.e., each of a doctor’s different offices.

²⁰ It does not capture, for example, sorting among medical facilities that might occur, for example, because of SARS. Facility-level visit levels (net of a patient’s visits) addresses this sorting.

our measure of peer effects $\Delta \bar{v}_{jkt}$ identifies only the effect of the change in average peer visits over the past year.²¹

We calculate the values of s_{jt} , \bar{v}_{jkt} , and $\Delta \bar{v}_{jkt}$ during the current week and the preceding 4 weeks and, for simplicity, sum up these values. This parsimonious approach measures the lagged effect of these variables on current visits, and is consistent with a potential delay between patient learning and patient action. Standard errors are clustered at the level of the patient's modal township, which allows for an arbitrary correlation across time periods for an individual, as well as among patients who visit networks in the same township.

The coefficient of interest in this specification is β_5 . If people respond to SARS by employing greater social learning, then this coefficient should be positive. In addition, we predict that the coefficients β_3 and β_4 will be negative because a higher underlying probability of catching SARS – as expressed in contact rates and SARS prevalence – should reduce demand for visits. We caveat our prediction about β_4 , the coefficient on total peer visits, however, because it may capture some peer effects. It is for this reason our analysis focuses on the change in average peer visits to cleanly identify social learning.

B. Data

Our empirical analysis is based on claims data from the Bureau of National Health Insurance (BNHI) in Taiwan. The BNHI provides fee-for-service health care to 96 percent of Taiwan's population. Claims data from this system provide a complete record of health care utilization since 1997 for a representative sample of one million patients. To economize on computing resources, we employ a random, one-percent subsample of this population

1. Peer groups and visits

We begin by defining a peer group to be the set of patients that visit a given doctor in a given clinic or hospital. Since a patient may visit several different doctors, we define several different peer groups for each patient. Since a doctor may visit several facilities, multiple peer groups may share the same doctor.²² A patient's set of peer groups is fixed during the period 1999-2000, before the earliest point in our primary data set. (Data from this period also provides information about the characteristics of patients and doctors.) To keep our description of data consistent with our theory, we will speak of doctors office and peer groups interchangeably.

Once we define all peer groups, we calculate the number of visits by each member patient and her peers to the doctor's office at the heart of the group. A patient's visits to a doctor's office are tabulated by week for the period from 2001 to 2003. Visits by a patient's peers are tabulated by

²¹ Moreover, \bar{v}_{jkt} and Σv_{jkt} capture nearly the same variation given that the number of patients visiting N_{jkt} each doctor's office is largely absorbed by doctor's office fixed effects.

²² The definition of social networks also excludes from the analysis visits to doctors that patients did not previously visit in 1999-2000, including doctors who had not begun to practice.

counting all visits to a doctor's office, again by week, and subtracting the patient's visits. Average visits by peers are normalized by the number of peers in the group as of the date the group is defined.

We construct our proxy for social learning by taking the average peer visits to a doctor's office in a given week by the average number of peer visits to that doctor's office in the same week the previous year. This is not the only approach for estimating the change in demand. One could, for example, subtract average peer visits three months prior or on a fixed date. But our approach does have two important features. First, it always compares the SARS period demand to pre-SARS demand since the SARS period is less than one year. Second, it controls for seasonality in health care attendance. Because we only use visit data for 2001-2003, only data from 2002 and 2003 are available for estimation after differencing.

According to our definition of peer groups, people belong to as many groups as doctor's offices that they visited during 1999-2000. Our regression analysis uses separate observations for each network to which a person belongs, but weights these observations by the fraction of a patient's total visits allocated to that network in 1999-2000. Thus patient observations from more frequently visited doctor's offices receive more weight than patient observations from less frequently visited doctor's offices. Moreover, each patient has equal weight in the overall analysis.

2. SARS cases

The Taiwan CDC provides SARS case data by date and location, distinguishing between "reported", "probable", and "confirmed" cases. We measure the number of SARS cases in a week by the number of reported cases that week, though all three measures are highly correlated and alternative measures do not affect our regression results. A person is classified as a reported case if he or she exhibits mild to moderate illness and meets the epidemiological criteria for possible exposure. Possible exposure is defined as travel to a location with documented or suspected recent transmission or recent close contact with someone else who was possibly exposed.

The probable case classification is more restrictive than the reported case classification. A person is classified as a probable case if he or she exhibits severe respiratory illness and likely exposure to SARS. Likely exposure is defined as close contact with a confirmed (in the laboratory) SARS case, or close contact with someone who is linked through a chain of transmission to a laboratory-confirmed SARS case.

Neither reported or probable cases are necessarily actual SARS cases. Scientists only developed the laboratory confirmation technique midway through the epidemic. Even then, confirmation of SARS was not available until several days after case reporting. Data on reported cases represent most closely the actual information available to patients at the time. A case typically received a "reported" designation immediately, at which time it was reported by the media. Public health officials subsequently updated the case designation to probable and then confirmed if merited by the patient's symptoms, lab tests, and epidemiological linkages.

3. The SARS epidemic period

The SARS epidemic may be subdivided into three distinct phases. The first phase, from March 10 to April 23, is characterized by intermittent SARS cases but no indigenous SARS transmission and no appreciable decline in health care attendance. The second phase, from April 24 to June 10, was the period of peak SARS transmission, during which time visits dropped precipitously. The third phase, from June 11 to September 23, is characterized by the diminution and disappearance of new SARS cases and the continued depression of health care attendance. The baseline regressions incorporate a single SARS period that is the union of phases 1 - 3 above. This window captures the period during which health care attendance was abnormally low. Subsequent analyses distinguish the degree of social learning during each of these phases and, moreover, during each week of the epidemic.

4. Summary statistics

Table 1 provides summary statistics for the key variables in our model – patient visits, average peer visits, total peer visits, reported SARS cases,²³ and year-on change in peer visits – for three time periods: pre-SARS, during SARS, and post-SARS. The data span calendar years 2002 and 2003 and the pre-SARS period runs from January 1, 2002, to March 9, 2003. The SARS period runs from March 10, 2003 – September 23, 2003. The post-SARS period runs from September 24 – December 31, 2003. Although the regression analysis will aggregate weekly data on regressors for five weeks and will weight multiple observations on patients so as to ensure each patient has equal weight, the statistics in Table 1 are neither aggregated nor weighted. The last column of the table provides p-values from a t-test comparing means of each regressor during the SARS period and outside the SARS period. These verify that there was a statistically significant drop in visits during SARS of roughly 20%.

Although Table 1 is self-explanatory, it omits data on aggregate demand for visits and the scope of peer groups. Taiwanese go to the doctor quite often and visit a variety of different doctor's offices. Eighty-four percent of Taiwanese visit the doctor at least once each year. Among this group, the median number of visits is 11, the interquartile range is 5 to 20, and the mean is 14.96. The median time between visits is only 4.7 weeks. Finally, the median patient visits 6 different doctor's offices and the interquartile range is 3 to 9. Conversely, the average doctor's office is visited by about 454 different patients.

IV. RESULTS

Our main regression results are reported in Table 2. The most parsimonious specification, which appears in Column 1, includes only our main treatment variables: the change in average peer visits $\Delta \bar{v}_{jkt}$ and that change interacted with the SARS epidemic indicator $S \Delta \bar{v}_{jkt}$. The latter identifies the effect of SARS on demand for social learning. Column 2 builds upon this specification by including total peer visits Σv_{jkt} and its interaction with the SAR epidemic indicator $S \Sigma v_{jkt}$ as controls. Column 3, which matches our full baseline specification (5), also includes the total number of reported SARS cases during

²³ Reported SARS cases (8.6/100,000 population) are higher than actual SARS cases (3.7/100,000) mentioned earlier in the paper because only a fraction of reported cases are ultimately confirmed as actual SARS cases.

the current week and the previous four weeks, as well as the interaction of this variable with the SARS period indicator.

The table shows a significant effect of $S\Delta\bar{v}_{jkt}$ and thus social learning on visits across all three specifications. In Column 1, a one unit increase in year-on change in average peer visits increases the patient's visits by 0.09 (se = 0.012). Controlling for additional variables attenuates this effect, but does not eliminate it. Column 3 reports that, conditional on total peer visits and SARS cases, the effect of $S\Delta\bar{v}_{jkt}$ remains significant with a magnitude of 0.068 (0.012).

Columns 2-3 also report the effects of total peer visits during the SARS period and of SARS cases that are consistent with our predictions. These two variables, which proxy contact rates and *inter alia* the probability of infection given contact, lower demand for visits by 0.04 (se = 0.01) and 0.042 (0.017).

Without more information, like the slope of demand curves, it is difficult to translate the coefficient on our measure of peers' beliefs into an economically meaningful value. However, it is possible to provide some context that will shed light on whether the coefficient on peer beliefs is economically significant. A one percent increase in our measure of peer beliefs (an additional change in average peer visits of 0.017 per week) increases a patient's visits by 0.05 percent.²⁴ By comparison, a one percent increase in our measure of contacts in the waiting room (additional average peer visits of 0.065) or in the probability of infection given a contact (additional reported SARS cases of 0.023) decreases a patient's visits by roughly 0.11 and 0.007 percent.²⁵ Thus social learning has roughly half the impact of other, objective measures of SARS risk.

1. Basic robustness tests

If there is positive serial correlation in a patient's visits over time but the lag of a patient's visits are omitted from the regression, then our estimate of the coefficient β_2 on $\Delta\bar{v}_{jkt}$ will be biased. The reason is that lagged visits $v_{ijk,t-52}$ is correlated with lagged average visits $\bar{v}_{jk,t-52}$, which lowers $\Delta\bar{v}_{jkt}$. There may also be bias in the coefficient β_5 on $S\Delta\bar{v}_{jkt}$ if serial correlation changes during the SARS period. Since the coefficient on Σv_{jkt} is positive we suspect the correlation between lagged patient visits and lagged peer visits is positive. Thus the bias in $\hat{\beta}_2$ is probably negative. In Column 4 of Table 2, we add $v_{ijk,t-52}$ to the baseline specification to correct omitted variable bias from first-order

²⁴ More precisely, we estimate that a 1% increase in our proxy for peer's beliefs about SARS risk increases visits over several periods. The sum of these increases is equivalent to a 0.5% change in one period of visits. When demand is roughly linear, which is a reasonable assumption for small changes in demand, the change in equilibrium quantity is proportional to the tax implied by SARS. When there is no ambiguity aversion, the tax from SARS is $T = -[\ln(1 - q)]/R \approx q/R$, where q is the estimate of SARS risk from peers. So the percent change in our proxy for beliefs, holding the slope of demand and R constant, is equal to the percent change in peer beliefs. When there is ambiguity aversion, however, the tax from SARS is $T + A$ when A is basically a constant (since q is very small). So a 1% change in $T + A$ is greater than a 1% change in peers beliefs about q . Thus, with ambiguity aversion, the actual elasticity of visits with respect to peer beliefs may be smaller than 0.05. Intuitively, some of the observed change in $\Delta\bar{v}_{jkt}$ is attributable to peers' ambiguity aversion rather than peers' assessments of q .

²⁵ The means used to calculate elasticities are estimated using the same weighting scheme used in the baseline regressions.

autocorrelation in patient visits. As suspected, $\hat{\beta}_2$ rises to 0.036. Our estimate of social effects, $\hat{\beta}_5$, falls but remains significant at 0.055 (se = 0.012).²⁶

Although our baseline regression uses the average visits of peers (along with network fixed effects which largely absorb the effect of peer group size) to proxy for the number of individuals encountered in the doctor's waiting room or supply constraints imposed by the doctor, we could also have used the total of peer visits to serve the same function. As indicated in column 1 of Table 3, substituting $\sum v_{jkt}$ for $\Delta \bar{v}_{jkt}$ does not materially alter our estimate of social learning. The coefficient on our measure of peer beliefs rises from 0.068 to 0.076 (se = 0.012).

The measure of peer beliefs in our baseline regression is the year-on change in average peer visits to a doctor's office. To test whether our estimate of social learning is sensitive to the time interval over which we difference demand, we employed alternative differencing periods. The one condition we impose is that the difference be larger than the length of the SARS period so that the baseline demand, i.e., average visits before, is measured prior to the SARS outbreak. We find that the length of the time interval does not affect our findings. Column 2 of Table 3 reports the results when we use, for example, a six-month difference in average peer demand. The coefficient on the difference in demand is 0.061 (se = 0.013), which is similar to the coefficient reported in Table 2.

In the baseline regression, we simply sum each of the primary independent variables, $\sum v_{jkt}$, $\Delta \bar{v}_{jkt}$ and S_{jt} , for the current and previous four weeks. This crude specification has the benefit of simplicity, but does assume equal coefficients on variables for the current period and for each lag. This turns out not to affect estimates of impulse response. We estimated the baseline regression using separate values for the current and lagged values of the primary independent variables. The sum of the coefficients on the current and lagged values for each of these variables was roughly the same as the analogue coefficient from the baseline specification.

In order to test how strongly our model capture peer effects, we conduct a falsification exercise using placebo treatments. We construct placebo treatments by taking each patient and randomly assigning her to another peer group within, say, her township.²⁷ Then we regress that patient's visits on local SARS cases and, instead of visits statistics for her actual peer group, the visit statistics for the randomly selected peer group. If our theory that a patient learns from her peer group is correct, we should find no effect from randomly selected peer groups, even in a close geographic area around the patient. We do this placebo treatment exercise with random assignment within different geographic areas: the patient's township, the patient's county, or all of Taiwan. The larger the pool from which we

²⁶ We can reject that the coefficient on lagged patient visits is one.

²⁷ Two point of clarification about this exercise. First, since each patient may go to many different doctor's offices and thus be a member of many different peer groups, what we actually do is assign each actual patient-peer group pair to a random other peer group. This random assignment does not preclude a patient being assigned to her actual peer group, though that is unlikely. Second, the random assignment takes place only once rather than each week, so each patient is assigned the full time sequence of visit statistics for another peer group.

choose a random peer group the less powerful our falsification exercise.²⁸ We find that the coefficient on our measure peer beliefs during SARS is small and insignificant at each level of randomization: -0.0072 (se = 0.0092) for township, 0.0001 (0.0107) for county, and -0.0041 (0.0128) for Taiwan. These results increase confidence in our theory that change in own peer demand – our measure of peer beliefs – affects patient demand for medical care.

2. Investigating the nature of peer effects

How much do patients learn about the risks of visiting one doctor’s office from their experiences with or information about other doctor’s offices that they visit? To test this, we regress a patient’s visit to one doctor’s office on not just the variables in the baseline regression, but also on variables that measure the behavior of peers at other doctor’s offices that the patient visits. The results are reported in column 3 of Table 3. We find a patient’s visit to a given doctor’s office is much more responsive to peers who visit that office than peers who visit other offices. For example, the coefficient on our measure of peer beliefs at the given doctor’s office has a coefficient of 0.071 (se = 0.013) as compared to a coefficient of 0.041 (0.021) on beliefs at other offices.

The baseline specification assumes peer effects are constant throughout the SARS period. However, the SARS period includes three distinct phases. We described these in section III.B.3. The third phase, from June 11 – September 23, is especially interesting because it includes the elimination of SARS cases, by July, and the extended depression of visits, through September. One hypothesis suggested by our theory is that social learning effect explains the prolonged suppression in visits following even after SARS cases disappeared. More precisely, social learning had a larger influence on patients’ posteriors about SARS risk once “objective data” on SARS cases vanished.

To explore this possibility, we modify the baseline regression by interacting average peer visits and the change in average peer visits with each three SARS phases separately. As hypothesized, the coefficient on change in average peer visits rises from 0.053 (se = 0.017) in phase 2 to 0.087 (se = 0.014) in phase 3. When we break the SARS period into two week intervals, we find this pattern again. Coefficients on our proxy for social learning rise as the number of SARS cases falls.²⁹ This is illustrated in Figure 3, which plots the coefficients on $\Delta \bar{v}_{jkt}$ interacted with indicators for each two-week interval in 2003.

One surprising finding from effort to track the evolution of social learning during SARS is that peer effects have an insignificant 0.018 (se = 0.023) role in phase 1, which covers the initial handful of cases. Perhaps there is a minimum salience that a new public health threat such as SARS must reach before people seek out more information on it.

²⁸ The opposite is true if we seek to use the placebo treatment exercise to check if our method of clustering standard errors by township correctly adjusts standard errors to account for serial correlation in visits. The results in the text suggest it does.

²⁹ In future research we will add additional data from the post-SARS period to see when the influence of peers subsides or whether it remains permanently higher.

Thus far we have assumed common treatment effects across townships. It is possible, however, that SARS visited different townships at different times. To determine whether individual townships' experiences with social learning are accurately conveyed by the aggregated results in Table 3, we created a township-level indicators that track whether a township, in a given week, had not yet experienced any SARS cases (pre-phase), was currently experiencing a SARS case (during-phase), or was not now but had in the past experienced a SARS case (post-phase). We interacted these three mutually exclusive phases with Σv_{jkt} and our measure of peer beliefs $\Delta \bar{v}_{jkt}$. The results suggest a pattern similar to the aggregate pattern across phases 1 – 3. The coefficient on our measure of peer beliefs is 0.02 and insignificant (se = 0.025) in the pre-phase, 0.053 and significant (se = 0.021) in during-phase, and 0.081 and significant (se = 0.02) in the post-phase.

In future research we will test the hypothesis that positive reports of SARS cases reduces the role of social learning in a second way. Instead of examining longitudinal changes in the role of social learning, we will examine cross-sectional changes. Specifically, we will interact Σv_{jkt} and our measure of peer beliefs $\Delta \bar{v}_{jkt}$ with a measure of how close the nearest SARS case is to the township in which the patient lives. This measure of closeness will be three mutually exclusive indicator variables that indicate whether the closest current SARS case is in the patient's township, in a township that is medium distance away, or a township that is a long distance away. (We will employ different measures of medium and long distance, including physical distance and overlap between social networks.)

REFERENCES

- Bandiera, Oriana, and Imran Rasul. 2006. Social Networks and Technology Adoption in Northern Mozambique. *Economic Journal*. 116: 869-902.
- Cheng, Tsung-Mei. 2003. Taiwan's New National Health Insurance Program: Genesis And Experience So Far. *Health Affairs*, 22(3): 61 – 76.
- Conley, Timothy, and Christopher R. Udry. 2008. Learning About a New Technology: Pineapple in Ghana.
- Deri 2005 [complete].
- Duflo, Esther, Pascaline Dupas, and Michael Kremer. 2009. Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya. Working paper.
- Einhorn, Hillel J., and Robin M. Hogarth. 1986. Decision Making Under Ambiguity. *Journal of Business*, 59(4), Part 2: S225-S250.
- Eliasz, Kfir, and Andrew Schotter. 2006. Paying for Confidence: An Experimental Study of Preferences over Beliefs. Working paper.
- Ellsberg, Daniel. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75(4): 643-669.
- Epstein, L. G. and J. Zhang. 2001. Subjective Probabilities on Subjectively Unambiguous Events. *Econometrica*, 69: 265-306.
- Foster, Andrew D., and Mark R. Rosenzweig. 1995. Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, 103(6): 1176-1209.
- Ghirardato, P., and M. Marinacci. 2002. Ambiguity Made Precise: A Comparative Foundation. *Journal of Economic Theory*, 102: 251-289.
- Hsieh, Ying-Hen, Cathy W.S. Chen, and Sze-Bi Hsu. 2004. SARS Outbreak, Taiwan, 2003. *Emerging Infectious Diseases*, 10(2): 201-2006.
- Institute of Sociology, Academia Sinica. May 2003. Special Survey on SARS.
- Karon, John, et al. 1990. HIV Prevalence Estimates and AIDS Case Projections for the United States: Report Based upon a Workshop. *MMWR* 39, No. 7, 110-2,117-9.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6): 1849-1892.

- Ko, Chih-Hung, Cheng-Fang Yen, Ju-Yu Yen, and Ming-Jen Yang. 2006. Psychosocial impact among the public of the severe acute respiratory syndrome epidemic in Taiwan. *Psychiatry and Clinical Neurosciences*, 60: 397–403.
- Kremer, Michael. 1996. Integrating Behavioral Choice into Epidemiological Models of AIDS. *Quarterly Journal of Economics*, Vol. 111, No. 2, pp. 549-573.
- Kremer, Michael, and Edward Miguel. 2007. The Illusion of Sustainability. *Quarterly Journal of Economics*. 1007-1065.
- Kreps, David M., and Evan L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, 46(1): 185-200.
- Kunreuther, Howard, Jacqueline Meszaros, Robin M. Hogarth, and Mark Spranca. 1995. Ambiguity and underwriter decision processes. *Journal of Economic Behavior and Organization*, 26: 337-352.
- Lu, Hung-Yi, and James E. Andrews. 2006. College Students' Perception of the Absolute Media Credibility about SARS-Related News during the SARS Outbreak in Taiwan. *China Media Research*. 2(2): 85-93.
- Lua, Tsung-Hsueh, Yiing-Jenq Choub, and Chien-Shian Liou. 2007. Impact of SARS on healthcare utilization by disease categories: Implications for delivery of healthcare services. *Health Policy*, 83: 375–381.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 74(6): 1447-1498.
- Manski, Charles F. 1993. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*, 60: 531–42.
- Moretti, Enrico. 2009. Social Learning and Peer Effects in Consumption: Evidence from Movie Sales. Working paper.
- Munshi, Kaivan, and Jacques Myaux. 2006. Social norms and the fertility transition. *Journal of Development Economics*, 80: 1– 38.
- Osmond, Dennis H. 2003. Epidemiology of HIV/AIDS in the United States. HIV InSite Knowledge Base Chapter. Accessed at <http://hivinsite.ucsf.edu/InSite?page=kb-01-03>.
- Oster, Emily, and Rebecca Thornton. 2009. Determinants of Technology Adoption: Private Value and Peer Effects in Menstrual Cup Take-Up. Working paper.
- Philipson, Tomas. 2000. Economic Epidemiology and Infectious Diseases. 1761-1799. A.J. Cuyler and J.P. Newhouse, eds. *Handbook of Health Economics* (Elsevier Science).
- Rao, Neel, Markus M. Mobius, and Tanya Rosenblat. 2007. Social Networks and Vaccination Decisions. Federal Reserve Bank of Boston, Working Paper No. 07-12.


Sacerdote, Bruce. 2001. Peer Effects with Random Assignment: Results for Dartmouth Roommates. *Quarterly Journal of Economics*. 681-704.

Valleroy, Linda, Duncan A. MacKellar, John M. Karon, et al. 2000. HIV Prevalence and Associated Risks in Young Men Who Have Sex With Men. *JAMA*, 284(2): 198-204.


Zimbelman, Mark F., and William S. Waller. 1999. An Experimental Investigation of Auditor-Auditee Interaction Under Ambiguity. *Journal of Accounting Research*, 37: 135-155.

FIGURES

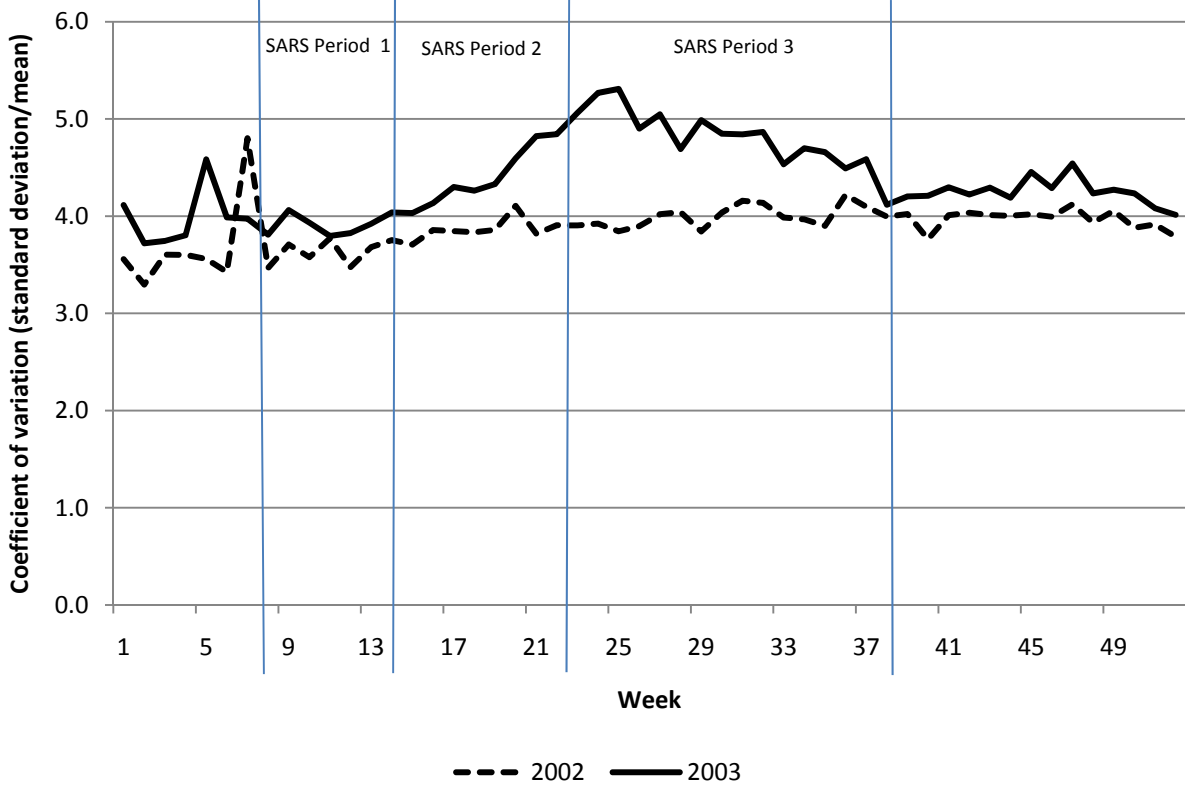
Figure 1. SARS pandemic in Taiwan.



B: Weekly Outpatient Visits During 2002 and 2003



C: Coefficient of Variation for Weekly Outpatient Visits During 2002 and 2003






Figure 2. Effect of SARS on equilibrium quantity is equivalent to tax on copay, where tax is proportional to the probability q of catching SARS.




Figure 3. Effect of the change in average peer visits on weekly visits during each two week interval in 2003.

TABLES

Table 1. Summary statistics by SARS period.

Time period:	Pre-SARS (1)	SARS (2)	Post-SARS (3)	P-Value (4)
Visits (per week)	0.0124 (0.1238)	0.0095 (0.1090)	0.0105 (0.1154)	0.00
Average visits by peers	0.0122 (0.0166)	0.0093 (0.0140)	0.0102 (0.0153)	0.00
Total visits by peers	5.5606 (7.8146)	4.2399 (6.3294)	4.6733 (6.9465)	0.00
Change in average visits by peers	-0.0038 (0.0146)	-0.0029 (0.0117)	-0.0016 (0.0118)	0.00
Reported SARS cases per 100 people	0.00 -	0.0086 (0.0037)	0.00 -	-
Sample size	3,601,928	1,712,392	826,672	

Notes. Standard deviations appear in parentheses. The sample is a 1% subsample during 2002-2003 of the BNHI panel of 1 million representative patients. The "Pre-SARS" period is from Jan 1, 2002, - Mar 9, 2003. The "SARS" period is from Mar 10 - Sept 23, 2003. The "Post-SARS" period is from Sept 24 - Dec 31, 2003. Column 4 displays the p-value from a t-test comparing the mean from Column 2 (SARS period) and the mean across Columns 1 (Pre-SARS) and 3 (Post-SARS). A peer group is all the patients who see a doctor at a given facility (doctor-facility pair). Peer groups are defined based on data from 1999-2000. "Visits" is the number of weekly visits to a doctor-facility pair by the index patient. "Average visits by peers" is the average number of weekly visits to a doctor-facility pair by members of the peer group, excluding the index patient. "Change in average visits by peers" is the difference in "Average visits of peers" from one year before. Observations are not weighted.

Table 2. Regressions of Weekly Visits on the Change in Visits of Peers.

	Dependent variable: visits (in week t)			
	(1)	(2)	(3)	(4)
Change in visits of peers (weeks t to t-4)	0.035 (0.008)	0.003 (0.008)	0.002 (0.008)	0.036 (0.009)
Visits of peers (t to t-4)		0.105 (0.018)	0.105 (0.018)	0.078 (0.018)
SARS x change in visits of peers (t to t-4)	0.090 (0.012)	0.068 (0.012)	0.068 (0.012)	0.055 (0.012)
SARS x visits of peers (t to t-4)		-0.040 (0.010)	-0.040 (0.010)	-0.038 (0.010)
Reported SARS cases (t to t-4)			-0.042 (0.017)	-0.036 (0.018)
Weekly visits (one year lag)				0.041 (0.004)
Patient fixed effects	Yes	Yes	Yes	Yes
Peer group fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Sample size	5,904,800	5,904,800	5,904,800	5,904,800

Notes. Standard errors, which appear in parentheses, are clustered by patient's modal township and are robust to heteroskedasticity. Regressions include all visits during 2002-2003 for a 1% subsample of the BNHI panel of 1 million representative patients. SARS is an indicator for the SARS Period (Weeks 10-38 of 2003). The remaining variables are defined consistently with Table 1. Where indicated, regressors are the sum of values for the contemporaneous week and four prior weeks. Regressions are weighted according to the share of the patient's visits that occur in the index peer group during 1999-2000.

Table 3. Regressions of weekly patient visits on visits of peers interacted with the SARS experience of patient's townships.

	Dependent variable: weekly visits to doctor's office j		
	total visits	6 month difference	other offices
	(1)	(2)	(3)
Change in visits of peers to office j (weeks t to t-4)	0.021 (0.008)	0.030 (0.009)	0.001 (0.007)
Change in visits of peers at other offices (t to t-4)			0.047 (0.015)
Visits of peers to office j (t to t-4)	0.000 (0.000)	0.078 (0.022)	0.102 (0.018)
Visits of peers to other offices (t to t-4)			-0.091 (0.019)
SARS x change in visits of peers to office j (t to t-4)	0.076 (0.012)	0.061 (0.013)	0.071 (0.013)
SARS x change in visits of peers to other offices (t to t-4)			0.041 (0.021)
SARS x visits of peers to office j (t to t-4)	0.000 (0.000)	-0.039 (0.009)	-0.038 (0.010)
SARS x visits of peers to other offices (t to t-4)			-0.010 (0.012)
Reported SARS cases (t to t-4)	-0.043 (0.017)	-0.038 (0.017)	-0.049 (0.017)
Sample size	5904800	5904800	5834200

Notes. Standard errors, which appear in parentheses, are clustered by patient's modal township and are robust to heteroskedasticity. Regressions include all visits during 2002-2003 for a 1% subsample of the BNHI panel of 1 million representative patients. SARS is an indicator for the SARS Period (Weeks 10-38 of 2003). Column 1 substitutes the total visits of the peer group for the average visits of peer group. Column 2 uses a 6 month lag to calculate the change in visits of the peer group, rather than a one-year lag in Table 2. Column 3 includes the visits of peer groups other than the index peer group. Where indicated, regressors are the sum of values for the contemporaneous week and four prior weeks. Regressions are weighted according to the share of the patient's visits that occur in the index peer group during 1999-2000.