

Scaling for Economists: Lessons from the Non-Adherence Problem in the Medical Literature

Omar Al-Ubaydli, John A. List, Danielle LoRe, and Dana Suskind

Many economists would be surprised to learn that patients adhere to the medications that physicians prescribe as little as 50 percent of the time (McDonald, Garg, and Haynes 2002). Clinical non-adherence is more than just an inconvenience to medical practitioners—it represents wasted resources and causes medical problems to evolve into forms that are even more expensive to treat. This has driven medical researchers to investigate rigorously ways of improving patient adherence. Their findings are of interest to economists who study interventions and wish to ensure that the inferences they draw from small-scale studies apply at larger scales, too.

More specifically, many experimental studies in economics are evaluations of a modification to an individual's behavior, based on the researchers' belief that such a modification will benefit the individual or confer benefits upon society. In the event that these beliefs are supported by the generated data, the broader goal is for large

■ *Omar Al-Ubaydli is a Researcher at the Bahrain Center for Strategic, International and Energy Studies, Manama, Bahrain, an Affiliated Associate Professor of Economics at George Mason University, Fairfax, Virginia, and an Affiliated Senior Research Fellow at the Mercatus Center, Arlington, Virginia. John List is Kenneth C. Griffin Distinguished Service Professor of Economics, University of Chicago, Chicago, Illinois, and a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. Danielle LoRe is a Medical Student at the University of Chicago Pritzker School of Medicine. Dana Suskind is Professor of Surgery and Pediatrics at the Duchossois Center for Advanced Medicine, University of Chicago, Chicago, Illinois. Al-Ubaydli is the corresponding author at omar@omar.ec.*

[†] For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.31.4.125>

doi=10.1257/jep.31.4.125

groups of individuals to then adopt the proposed intervention autonomously, or to have policymakers promote its adoption via methods such as subsidies, awareness campaigns, legislation, and so on.

For example, in an effort to provide policymakers with recommendations on how to improve agricultural productivity, Duflo, Kremer, and Robinson (2011) investigate the benefits accrued to Kenyan farmers who invest in fertilizer. Likewise, Field and Pande (2008) study the effect of loan repayment frequency on client default in microfinance by experimentally manipulating repayment schedules, with an eye to supplying creditors with scientific information about optimal loan structure.

These studies typically result in a recommendation to policymakers about how to affect broad-based behaviors outside of the experiment. The studies also show why people were deviating from the optimal behavior prior to study. For example, perhaps they simply did not realize that they could do better by modifying their behavior, as in the Kenyan farmer case (Duflo, Kremer, and Robinson 2011). Or infeasibility could be the driver, as in the case of studies investigating novel repayment schedules, or it could be something beyond the agent's financial means, such as Fryer's (2011) study of financial incentives in schools.

While scientists typically have clear advice for policymakers, a common occurrence is that such research programs are never scaled, or when they are scaled, the size of the measured treatment effect diminishes substantially relative to the that found in the original study. This is a common phenomenon known in the literature as "voltage drop," but this type of predictable change is not accounted for in benefit–cost analysis. The papers cited above have not, to the best of our knowledge, exhibited voltage drop—we mentioned them because they are archetypal experimental papers where voltage drop is a concern.

In terms of actual examples of voltage drop, consider significant public health threats, such as HIV, tuberculosis, and malaria. Despite the demonstrated effectiveness of drug therapies on transmission in trials and in small-scale settings, prevalence in developing countries remains high. For example, Global Fund to Fight AIDS, Tuberculosis, and Malaria sought to raise and disburse money to poor countries in an effort to provide therapies (Lu, Khan, and Murray 2006). However, half the funds never reached clinics due to inability of health services to manage the funds effectively (Garrett 2007), resulting in significantly weaker "effective" treatment effects. This voltage drop was caused by financial and resource constraints that were never overcome.

A related example includes initiatives to decrease rates of transmission and promote safe sex practice, such as providing condoms to a community. Such practices have faced voltage drops when scaling, due to variations in community beliefs and values (Campbell and Mzaidume 2002). In a review of barriers to HIV intervention implementation, stigmatism of prevention and treatment, power dynamics within society, and plateauing of health education messaging were all identified as decreasing efficacy of these provisions (Chopra and Ford 2005).

The experimental literature is littered with such examples. In a general sense, the issue revolves around a query of this form: "I just found a 0.2 standard deviation

effect in my experiment. Should I expect to observe such an effect when scaled to a city, state, or country?” Voltage drop can occur for many reasons, and it constitutes an example of a particularly vexing public policy challenge that we denote “the scaling problem.” Threats to scalability can be divided into three areas: statistical inference; representativeness of the experimental population; and representativeness of the experimental situation (Al-Ubaydli, List, and Suskind 2017).

The statistical inference class of scaling problem relates to inferential errors by scholars and policymakers seeking to apply insights gleaned from a small-scale study to a larger-scale setting. These errors typically relate to a failure to adjust correctly for the fact that a published statistical finding about the relationship between two variables is merely one of numerous, contemporaneous, investigations into the same relationship conducted by other research teams working quasi-independently (Maniadis, Tufano, and List 2014). This inference problem is related to the broader problem of publication bias (Young, Ioannidis, and Al-Ubaydli 2008).

The “representativeness of the experimental population” class of scaling problem refers to the challenge of ensuring that the subject pool in a small-scale study is representative of the larger population targeted by policymakers seeking to scale the findings. For example, does an experimental investigation of a proposed intervention that improves saving habits, which is conducted on college students, yield results that apply to the entire population, which is mostly composed of people who are not currently enrolled in college? Note that this class of scaling problem is not specific to experimental data—it applies to naturally occurring data, too. However, there exists a lively debate about the susceptibility of various data types to this type of scaling problem (Al-Ubaydli and List 2015b; Deaton and Cartwright 2016).

While both statistical inference and representativeness of the population are important, the focus of our discussion is on representativeness of the experimental situation. The experimental situation is quite rich and includes many important considerations. For instance, the next simple example of “program drift” illustrates one set of reasons for the scalability problem within the experimental situation. Consider Early Head Start home visiting services, one of the largest federally funded early childhood interventions in the world. The program demonstrated significantly improved school readiness for children aged up to three years old, improved family economic self-sufficiency, and parenting practices through high-quality efficacy trials (Paulsell, Avellar, Martin, and Del Grosso 2010). However, variation in quality of home visits was found at larger scale, with home visits for “at risk” families involving more distractions and less time on child-focused activities, causing the delivery of a different program than what had been studied. Lower proportion of time on child-focused activities and lower parental engagement was associated with diminished effectiveness for both child and parent outcomes as well as higher dropout rates (Raikes, Green, Atwater, Kisker, and Constantine 2006; Roggman, Cook, Peterson, and Raikes 2008). General equilibrium effects also fall under this class of scaling problem, whereby fidelity to the original small-scale design at a larger-scale setting results in interactions with other variables, in turn causing treatment effects to structurally change.

Closely related to delivery of the wrong program, or the wrong dosage of the program, is that no program at all is received. For instance, individual non-adoption of treatment represents a serious consideration when assessing the efficacy of public programs. Even when non-adoption apparently contradicts the best interests of agents, it is often found within government programs or after findings from research studies are made public. For example, many scholars are puzzled by the persistent reluctance of consumers to purchase energy-efficient lightbulbs despite the manifest cost savings that they offer and the absence of any notable downside to their usage (Allcott and Taubinsky 2015). Likewise, when new technologies are advanced as governmental policies, such as agricultural, financial, or time-saving computer technologies (like renewing one's passport online), the level of adoption is typically far less than most anticipate.

Medical practitioners have for centuries been facing an isomorphic problem: patient non-adherence to prescribed medications, specifically those that manifestly serve the patient's interests as opposed to cutting-edge, experimental medicines without established efficacy or medicines that require invasive administration with high risk of complications or intolerable side effects. This problem has spawned a large literature in the medical sciences regarding the best practices for improving medication adherence, and the results are directly relevant to economists seeking to tackle the narrow component of the scaling problem considered in this paper. In this study, we explore the findings of medical practitioners and present a series of recommendations tailored to the environments typically studied by economists.

One reasonable response to the problem of non-adherence when scaling is to interpret this as a failure to design the original, small-scale study properly. Thus, rather than proposing ways to enhance adherence, we could focus on how to design the original experiment such that it correctly captures the expected level of adherence in the larger scale. We regard the two approaches as complementary. We choose to focus on techniques for boosting adherence because low adherence in the general population undermines treatment effects and therefore the effectiveness of policy. Moreover, empirical researchers often gather experimental data precisely because the enhanced control possible in the experimental setting allows for cheaper and more powerful estimates of treatment effects, as a precursor to more effective policy. In other words, the low adherence levels observed in large-scale field settings should not be taken as an inescapable constraint; part of the study should involve considerations to boost adherence.

Scaling and Researcher Control over Non-Adoption in Economics Experiments

Controlled experiments, be they laboratory or field, have become mainstream in economics only during the last 30 years, and therefore scaling issues, including the specific one under consideration in this study, are relatively recent problems for economists. Economics experiments typically involve unusual levels of control over

the options available to agents. In the case of laboratory experiments as well as most field experiments (with the exception of natural field experiments), researcher control is close to absolute, allowing scholars to severely restrict choice sets: for example, *cooperate* or *defect* in a prisoner's dilemma game (Andreoni and Miller 1993), or *get paid to go to the gym* or *do not get paid* (Charness and Gneezy 2009).

Even in the case of natural field experiments, researchers often select natural environments that offer enhanced levels of control. For example, Shearer (2004) experimentally compared piece-rate and flat compensation schemes in a natural setting by finding a rare case of a company that uses both compensation schemes in its operations for the same type of work, and would therefore be willing to experimentally (and covertly) modify the compensation scheme offered to its workers.

Elevated levels of control are a key reason for the attractiveness of experimental methods when seeking to evaluate the consequences of a proposed modification to agents' behavior. For example, in the Duflo, Kremer, and Robinson (2011) fertilizer experiment, a first step toward demonstrating the superiority of fertilizer, and therefore convincing farmers of the benefits of using fertilizer, is heavily subsidizing the fertilizer or the delivery thereof. Similarly, Gosnell, List, and Metcalf (2016) were able to alert airline pilots and their managers to the possibility of enhancing fuel efficiency by offering pilots elevated levels of feedback, and by providing them the opportunity to earn money for charity.

Many of these evaluative studies result in a conclusion of the form: "upon being made aware of our findings, agents should autonomously adopt the behavior that was investigated," based on the premise that a primary barrier to the previous adoption of the behavior was informational. Yet in practice, economists often find lukewarm receptions for their findings among the agents who should respond by updating their behavior—whether the audience is policymakers, firms, or laypeople. When Ferraro and Price (2013) demonstrated that providing social comparisons in utility bills yields significant improvements in conservation at a trivial cost, utility companies the world over should have, in principle, expressed interest in deploying similar policies, yet this has not occurred. Likewise, Hossain and List (2012) discovered that the productivity-enhancing effects of providing Chinese factory workers with financial incentives led to a net increase in profits, and that this effect was larger when the incentives were presented in a negative frame—the most novel component of the experiment. According to neoclassical economics, merely publicizing this finding should lead to substantial enthusiasm for the adoption of such methods, and the scaling of the result. To our knowledge, this has not occurred widely.

Admittedly, in Ferraro and Price (2013) and Hossain and List (2012), and more generally in the case of the thousands of other economics experiments conducted, some of the reluctance among agents to modify behavior is due to uncertainty over the generalizability of the finding in question—what works for a Georgia water company might not work for a Slovenian electricity provider. It may work in China, but does it work in Toledo?

However, we can be confident that part of the non-adoption can be classified as purely irrational behavior, as many of the relevant studies investigate the promotion of actions that people should already be undertaking themselves. This includes Charness and Gneezy's (2009) use of financial incentives to induce greater exercise and Allcott and Taubinsky's (2015) attempts at increasing usage of energy-efficient lightbulbs. Moreover, in laboratory and field experiments (except natural field experiments), purely irrational behavior may be temporarily suppressed by experimenter demand effects or by the artificial restrictions on choices available to participants (Levitt and List 2007), accentuating the discrepancy between agents' willingness to modify their behavior in the study and in the natural environments ultimately targeted by the researchers.

This state of affairs poses a problem for policymakers seeking to scale empirical findings. If policymakers rely purely on publication of the results, then adoption will be impaired by irrational non-adoption. Alternatively, should the policymakers try to replicate the methods used in the original study, then they will face a host of structural scaling problems, such as the rising marginal cost of program administration, heterogeneity in the population, intransigence by stakeholders who are invested in the prevailing mode of behavior, and a litany of other issues discussed more fully in the rest of this symposium and in Al-Ubaydli, List, and Suskind (2017).

To illustrate these issues with a concrete irrational non-adherence example from the economics literature, consider the small-scale study conducted by Fryer, Levitt, and List (2015). Using a sample of 257 families from Chicago, the authors studied the effect of providing parents with financial incentives to engage in behaviors designed to increase early childhood skills via a parent academy that delivered training sessions. The study found large and statistically significant effects; in particular, over 80 percent of parents attended at least one training session, and over 40 percent attended all sessions, which is a crucial link in the causal chain under investigation. Inspired by these findings, the UK Education Endowment Foundation launched a parenting academy and a study structured similar to that in Chicago, but involving over 2,500 children spread across a larger geographical area. The larger-scale program found that only 60 percent of parents attended at least one session, and only 11 percent attended all sessions. Unsurprisingly, with such weak attendance, the study found no evidence of a positive effect of the interventions (in fact, the absence of an effect was true even when controlling for attendance).

Examples as clear as this are rare in economics, simply because this system of small-scale experimental research leading to large-scale policy implementation is a recent addition to the discipline. We anticipate that such problems will increase in frequency over the coming years as a larger volume of the profession's research resources are dedicated to this system for delivering policy insights. In this spirit, we envision nonprofit and for-profit firms, governmental bodies from local to federal, and supranational authorities as strong demanders of information on the causal effects of interventions.

If this is indeed the case, then there are considerable benefits associated with devising methods to deal with irrational non-adoption, and the first step is to obtain

a better understanding of the underlying causes. The large behavioral economics literature offers many convincing explanations; based on our experience, we draw attention to the following likely sources.

First, humans experience *psychological switching costs* (Klemperer 1987; Carroll, Overland, and Weil 2000) and tend to dislike modifying their behavior for reasons independent of any material cost associated with changes in behavior. This may be a manifestation of an overarching tendency for humans to exhibit path dependence, which commonly surfaces in the form of the endowment effect (Kahnemann, Knetsch, and Thaler 1991). It may also reflect a propensity to herd: that is, to avoid deviating from the manner in which peers are behaving (Chang, Cheng, and Khorana 2000). In the context of scaling small-scale experimental results, psychological switching costs and habit formation constitute a barrier to the organic modification of behavior prescribed by a study.

Second, humans often display *hyperbolic discounting* (Laibson 1997), that is, when the cost is borne up front, they can indefinitely delay decisions that serve their interests because of an irrational fixation on reaping short-term rewards. This model of decision-making is used to explain apparently irrational patterns of credit card usage, such as borrowing at a high interest rate while simultaneously depositing money in a checking account (Telyukova and Wright 2008), as well as irrationality in pensions and savings decisions (Thaler and Benartzi 2004). Thus, even when a small-scale experiment demonstrates the benefits of a modifying behavior, agents may still exhibit reluctance toward organically adopting the change if it requires an up-front cost, despite the back-loaded benefits more than offsetting the up-front costs.

Third, when *complexity is combined with limited cognitive capabilities* (Simon 1972), humans may sometimes wish to take a certain course of action in the pursuit of their interests but be prevented from correctly modifying their behavior by limited cognitive abilities. For example, many who succumb to the Allais (1990) paradox (which involves choices between different sets of gambles) lack the intellectual capacity to understand the potential sub-optimality of their actions. In the context of insurance, consumers exhibit significant difficulty in making rational assessments of the premiums and deductibles offered in contracts (Watt, Vazquez, and Moreno 2001). Similarly, people may have systemically incorrect beliefs about the consequences of actions (Caplan 2002). In the context of small-scale experiments, not all agents are equipped with the cognitive tools necessary to appreciate the benefit of a prescribed change in behavior, or to acquire and process the information required to make a sound judgment.

These explanations are not intended to be exhaustive; our aim is simply to illustrate that the economics literature provides us with rich refinements to the baseline neoclassical model of decision-making that can account for why people sometimes seemingly refuse to pick up the proverbial dollar bills from the sidewalk. While economists have investigated a broad range of appropriate countermeasures, when it comes to the problem of getting people to modify their behavior for their own

benefit, we can draw important lessons from attempts to address one version of this problem in medicine.

An Isomorphic Problem: Medication Non-Adherence

Clinicians take great care to ensure that the medications they prescribe to their patients serve their patients' interests. This includes social norms such as the Hippocratic oath (Orr, Pang, Pellegrino, and Siegler 1997), a sophisticated system of oversight by clinical peers and administrators (Farnan et al. 2012), and the threat of legal action in the event that clinicians fail to serve patients' best interests (Studdert et al. 2006). Moreover, clinical units that underperform suffer adverse commercial consequences, as consumers care about reputation (Hibbard, Stockard, and Tusler 2005).

While no system is perfect, medical practitioners should be considered highly motivated to provide sound prescriptions that patients can trust. Despite these favorable conditions from the perspective of patients, typical adherence rates for prescribed medications are around 50 percent (McDonald, Garg, and Haynes 2002), a figure that is of grave concern for clinicians because it diminishes the benefits of the treatments. This non-adherence rate has broad implications, including raising the costs of healthcare, prolonging patient discomfort, allowing disease progression, and biasing assessments of the effectiveness of treatments (Vervloet et al. 2012).

What drives so many patients to behave consistently in a manner that apparently contradicts their best interests?

Causes of Medication Non-Adherence

The medical literature has identified two primary classes of cause for medication non-adherence: intentional and unintentional (Kripalani et al. 2007; Vervloet et al. 2012; Dayer, Heldenbrand, Anderson, Gubbins, and Martin 2013). Interestingly, these two causes overlap with the causes pinpointed for irrational non-adoption discussed in the economics literature described above.

Intentional non-adherence, which refers to willful cost-benefit analysis by the patient, usually results from the patient attaching significant discomfort to the medication, and assessing that the purported benefits from the medication do not justify the discomfort. For example, a liquid medicine may have a disagreeable taste or may need to be administered via a painful injection. Intentional non-adherence may be exacerbated by systematically inaccurate beliefs on the effects of a medication, such as when patients prefer anecdotal evidence, or the advice of a celebrity, to the results of formal studies.

One such example helps to illustrate this mechanism at work—recall that the decline in vaccination rates and concurrent rise in vaccine-preventable diseases align with an anti-vaccine movement that, despite significant scientific evidence to the contrary, was fueled by personal stories and celebrity endorsement. At the time,

such information outweighed pro-vaccine information on user-friendly outlets such as YouTube (Venkatraman, Garg, and Kumar 2015).

There are, of course, situations where intentional non-adherence can be a rational personal choice. In terminal disease and end-of-life care, it is straightforward to make the case that the small gain in lifespan provided by a medication is objectively not worth the decrease in quality of life. However, for most situations, such non-adherence corresponds to economists' model of hyperbolic discounting, possibly combined with cognitive limitations and/or biased beliefs.

Unintentional non-adherence is a major hurdle for clinicians and patients. It covers simple forgetfulness, as well as failures to comply with treatment plans resulting from regimen complexity, which can stem from quantity of medications and frequent dosage times or complicated, multistep administration of medicine (as with inhalers); as a result, treatment may be carried out incorrectly and thus ineffectively (Lavorini et al. 2008). Physical problems, such as sleeping through a scheduled treatment appointment due to fatigue, or lacking the mobility to adhere to a regimen, are also included in this class of cause.

Clinicians have devised a diverse range of interventions to address these factors, and have tested them using randomized control trials. Before evaluating these interventions, it is worth considering what can be inferred about medication adherence from naturally occurring variation in treatment features and background variables. Summarizing the literature broadly, McDonald et al. (2002) conclude that compliance is at best weakly related to sociodemographic factors, including age, sex, race, intelligence, and education.

Interestingly, McDonald et al. (2002) also found that patients with physical disabilities caused by the disease being treated were more likely to adhere to the prescribed regimen. Clearly, it is difficult to ascertain the degree to which such results generalize to the domains most frequently encountered by economists, but this finding may reflect simple cost–benefit dynamics: that is, those who benefit most from a treatment make the most effort to adhere. For example, Alcott and Taubinsky (2015) report that a significant proportion of non-adoption of energy-efficient lightbulbs might potentially be attributable to the fact that the financial returns of switching—while being positive in net terms—were too small to justify the act. Furthermore, and as expected, natural increases in the cost/complexity/duration of treatment plans are also associated with diminished compliance.

Improving Medication Adherence: Methods

Improving medication adherence is a central problem in the medical sciences, spawning thousands of papers and dozens of meta-studies. Peterson, Takiya, and Finley (2003) provide a useful categorization of the types of interventions that practitioners have evaluated in formal trials, many of which should be instantly recognizable to experimental economists who conduct small-scale studies with the goal of scaling their results to larger populations.

One important class of studies is *educational interventions*, whereby the medical team attempts to plug any informational lacunae that the patient may be suffering.

For instance, they may provide instructions on how to take a specific medicine and address misconceptions that the patient might have regarding the treatment's effectiveness or its side effects. Educational interventions vary along many dimensions, such as the medium (oral, visual, written), the delivery method (in person, telephone, electronic, printed), the professional delivering the intervention (physician, nurse, pharmacist), the frequency (one-off, weekly, monthly), the location (home, hospital, community center, remote), and the number of participants (one-to-one, group).

Related to educational interventions are *counseling and accountability interventions*, whereby members of the clinical team follow-up with the patient on the treatment to ensure that it is being taken as prescribed. In addition, under this approach, informational and psychological support are provided as the need arises. These interventions may feature monitoring devices, such as remote blood pressure sensors, that assist clinicians in gathering accurate information about a patient's adherence to best provide personalized counseling.

Medical practitioners have also investigated *interventions that support the patient's independent adherence efforts*, such as self-monitoring devices, including simple pill-boxes that help patients track how many pills they have ingested, as well as more sophisticated electronic aids that measure vital signs. The advent of mobile telephone technology has greatly enhanced the opportunity to make use of automatic reminders, including text messages and notifications from smartphone applications.

An intermediate form of intervention, familiar to economists studying micro-finance, is *involving family members in counseling and educational sessions*. Family members can directly assist in the delivery of treatment (for example, by injecting a patient who might otherwise be reluctant to inject themselves), provide reminders and emotional support, and give clinicians richer feedback on the degree of adherence and on the sources of non-adherence.

A final intervention class that—to the best of our knowledge is scarcely deployed in the medical non-adherence literature—is *using financial incentives*. While clinicians regularly advocate subsidizing treatment plans up to the point of free provision, reflecting a tacit acceptance of the importance of financial considerations in the patient's adherence calculus, there appears to be very little appetite for actually *paying* patients to take medicines as prescribed. This holds even if a reasonable cost-benefit case can be made in terms of the medical authorities avoiding more expensive treatments further down the road arising from non-adherence at present (Guiffrida and Torgerson 1997).

Several reasons have been suggested for this comparative rarity of financial incentives for medical adherence. For example, one concern may be due to clinician awareness of the debate regarding extrinsic versus intrinsic incentives (Deci, Koestner, and Ryan 1999; Benabou and Tirole 2003) and the fear that extrinsic incentives may diminish intrinsic motivation. Alternatively, there may be fears that positive financial incentives could induce spurious claims for the need for treatment. Yet another possibility based on sunk cost reasoning could account for

clinician reluctance: if the individual pays a positive amount for medication, then they are more likely to use the product than if they receive it for free.

A field experiment due to Ashraf, Berry, and Shapiro (2010) suggests that this might be the case for households using Clorin to treat drinking water: although they find no evidence that people who paid lower prices consume the product less than those paying higher amounts, they find some evidence suggesting that those who pay nothing use it least. Notably, some clinicians have successfully used financial incentives to encourage general healthy behavior, such as smoking cessation (Halpern et al. 2015) and weight loss (Volpp et al. 2008).

The diversity in methods adopted by clinicians to address non-adherence is partially a response to the diversity of conditions treated, and consequently the diversity of treatments. In our experience, economists sometimes view clinical medical trials as a binary situation: take the drug or don't. But in fact, medical conditions vary in a large number of dimensions: whether they are one-time or chronic, the nature of the discomfort that they induce, the time profile of the condition's effects upon the patient, the efficacy of the treatments, the results of noncompliance, and so on. Of course, economic environments feature parallel levels of context-specificity. Our point here is that when economists consider the medical literature on non-adherence, they should be aware that issues of context arise here, too.

Improving Medication Adherence: Traditional Results

What have clinicians learned and surmised based on randomized control trials designed to improve medication adherence? We will focus in this section on what we call "traditional" meta-studies, which covers most meta-studies conducted up to around 2012. These studies exclude investigations of smartphone applications and other mobile telephone-based methods of improving medication adherence, which represent more recent technological innovations. In the next section, we focus on the effectiveness of modern mobile telephones in clinicians' quest to enhance medication adherence. Because our ultimate focus is applying these results to the environments that typically interest economists, which are quite distinct from those considered in the medication non-adherence literature, we focus on providing readers with qualitative results. Those interested in a quantitatively rigorous meta-analysis should consult the meta-studies cited here. We primarily draw upon the work of McDonald, Garg, and Haynes (2002), Peterson, Takiya, and Finley (2003), Kripalani, Yao, and Haynes (2007), Haynes, Ackloo, Sahota, McDonald, and Yao (2008), Zullig, Peterson, and Bosworth (2013), and Nieuwlaat et al. (2014).

An overarching—and somewhat disappointing—conclusion from this literature is that the methods considered exhibit a high degree of context-specificity in their effectiveness, making it difficult to arrive at general conclusions. As mentioned above, this is the result of the huge diversity in medical conditions, and in the treatments that clinicians prescribe in the pursuit of better health outcomes. Consequently, this sobering conclusion should not be considered anomalous. Interestingly, such results parallel the arguments in Levitt and List (2007) concerning the generalizability of experimental results from the lab.

Overall, a slight majority of studies find no significant effect of the interventions being investigated on medication adherence. In fact, it is quite common for some very expensive interventions, such as face-to-face meetings with specialist physicians, to result in no statistically discernable effect upon medication adherence. Yet, it is important to note that experimental power may be a culprit here—most studies are not reporting a treatment effect estimate of precisely zero.

Among the approximately 40 percent of studies that do detect a statistically significant effect, the magnitude is somewhat modest, falling in the range of 4–11 percent. Importantly, detected effects tend to shrink further when one focuses on the relationship between the adherence intervention and clinical outcome, rather than the intermediate relationship between the adherence intervention and the rate of adherence. Moreover, there is no general pattern regarding the comparative effectiveness of narrow interventions, such as focus groups versus email reminders.

Particularly in the case of long-term, chronic medical conditions however, there is a tendency for the most effective interventions to be those based on complex combinations of the basic classes, such as educational sessions at the start, counseling sessions throughout the treatment plan, and a selection of reminder methods, such as telephone calls from nurses and pillboxes.

Finally, we should highlight that the value of studies of medication adherence is limited by a series of flaws in the data-gathering and analysis process. First, datasets tend to be small and experimental designs underpowered, and authors of survey articles typically urge scholars to pay more attention to established best practices in sample size determination. Second, the somewhat inevitable dependence upon self-reported measures of medication adherence, especially in the traditional studies that predate the era of smartphones and remote monitoring, is a considerable source of noise that impedes precise inference. Third, publication bias—the tendency for journal editors to systematically favor studies that report significant results—is a source of upward bias in detected treatment effects, though (as with issues of experimental design), appropriate coordination between scholars and journal editors can eliminate this problem (Young, Ionnidis, and Al-Ubaydli 2008).

Improving Medication Adherence: Smartphone Results

The traditional literature suggests that organic medication adherence rates can be quite modest, and that exogenous interventions tend to have a small effect at best on patient adherence to prescribed medications. Given the dramatic effects that smartphones have had on the nature of many services delivered to consumers, such as banking, dating, ridesharing, and media, there is a sense of optimism that they can also contribute to higher rates of medication adherence.

In particular, smartphone applications have several novel and attractive attributes (Dayer et al. 2013): constant accessibility, the ability to act as a repository of patient- and medication-specific information; a source of education for patients about adherence; and interoperability with existing systems for prescriptions and medical records. Critically, the cost of these features is potentially many orders lower than that of the next-best alternative. For example, if a patient has to convey

self-reported adherence to a clinician in oral or written form, there is a considerable time cost of recording the data in the patient's medical file, as compared to the instant integration that a smartphone can offer.

It is too early for the literature to provide rigorous quantitative assessments of the effectiveness of smartphone interventions targeting medication adherence. Existing studies, such as Dayer et al. (2013), focus more on qualitative conjectures about the likely effectiveness of various features. They emphasize the positive role of several features, including the ability to sync adherence data with records housed on servers of healthcare providers. In addition, the ability to track missed and taken doses, not just via patient self-reporting, but also via direct cable or wireless link to various treatments can be important. Further, the ability to provide detailed instructions on complex medications—the value of which can be enhanced by linking to databases that give patients a broader background on their medical conditions and treatments—is invaluable. Finally, such features can address multilingualism, as the technical vocabulary associated with clinical settings can be a challenge for the millions of migrants that live in a country with a language different than their mother tongue.

Another piece of evidence lending insights into the potential efficacy of smartphone technology is Vervloet et al. (2012). They report that, across several studies, electronic reminder devices—which operate in a manner similar to pager systems, allowing for automated or visual reminders—have a substantial, positive, and robust effect on medication adherence. This is interesting evidence because it serves to highlight the potential role that smartphones can play, seeing as electronic reminder devices are in many regards rendered obsolete by smartphones, which can perform all of the same functions, as well as many additional ones, described above.

As an illustration, one smartphone application formally evaluated in a randomized control trial was the WellDoc diabetes management application, which displays medication regimen, provides feedback on patients' blood glucose levels, and tailored management through evidence-based algorithms (Quinn et al. 2008). Compared to a control group, patients showed a significant decrease in HbA1c, a clinical measure of diabetes management, and its success was salient enough to convince insurance companies to subsidize the application, allowing it to be prescribed as part of the treatment.

Nevertheless, the overwhelming majority of quantitative studies of smartphones, such as Strandbygaard, Thomsen, and Backer (2010), Petrie, Perry, Broadbent, and Weinman (2012), Huang et al. (2013), and Finistis, Pellowski, and Johnson (2014), are forced to focus on the simplest intervention that smartphones permit—text message reminders. The comprehensive surveys by Vervloet et al. (2012) and Sarabi, Sadoughi, Orak, and Bahaadinbeigy (2016) concluded that there is robust evidence that text message reminders improve medication adherence, especially in chronic conditions or in patient populations requiring complex medication regimens, such as HIV, asthma, and diabetes. There is evidence that tailoring the messages to the patient yields a larger effect on medication adherence (Kreuter, Farrell, Olevitch,

and Brennan 2000), as do interactive messages—for example, ones that require a reply from the patient.

An additional finding from the literature is that reminders are systematically more effective for those who are unintentionally non-adherent, such as older patients who have memory problems, or adolescents who might be preoccupied with their social lives. In contrast, reminders are found to be ineffective for the intentionally non-adherent, which further illustrates the need for smartphone capabilities that go beyond reminders to most fully address non-adherence. Many patients have an increasing interest in accessing health information on smartphone apps over internet sources (Smith 2015), which provides a promising avenue for dissemination of scientific evidence in a user-friendly manner to combat intentional non-adherence due to systematically inaccurate beliefs.

As a whole, the consensus regarding smartphones at this point is that there are sound reasons for optimism, but there remains a need for the accumulation of further evidence. There is the possibility that in the long-term, such interactivity through smartphones might backfire by creating user fatigue, especially in light of the variety of stimuli that smartphones offer (Dennison, Morrison, Conway, and Yardley 2013). One possibility is that future consumers might use a number of wearable electronic devices, perhaps including a device dedicated to medications or to fitness more broadly, rather than bundling so much of their personal information technology into a smartphone. In fact, in clinical settings, the risks posed by smartphones, including their acting as a distraction for patients and clinicians, have driven some medical researchers to propose strict regulations on smartphone usage in conjunction with using specialized alternatives, such as electronic reminder devices (Gill, Kamath, and Gill 2012).

What Can Economists Learn from Medicine?

In the narrow domain of irrational non-adoption, there is much that economists can potentially learn from medical researchers, as the latter group has been rigorously studying an isomorphic problem for decades. In addition, these studies are in a setting where the stakes are significantly higher than those typically encountered by economists. Several insights can be gained from the medical adherence literature.

First, economists should not assume that merely demonstrating the superiority of an alternative mode of behavior to an agent—even from the perspective of the agent's interests—is sufficient for the agent to organically modify their behavior. Admittedly, almost all economists understand this point, as confirmed by the existence of a large and heavily cited behavioral economics literature. However, our sense is that this lesson is sometimes forgotten when economists are solicited by policymakers interested in scaling up their observed findings. Maybe it is due to the understandable excitement of seeing one's research have a profound effect on society—a rare event in the life of most economists. Alternatively, it could be

the result of the difficulty of explaining such subtleties to the nonspecialist policymakers and senior civil servants who expressed interest in scaling the findings. Yet exercising restraint is critical at such junctures.

Clinicians have struggled with patient medication adherence rates that average around 50 percent for decades, due to reasons ranging from willful noncompliance stemming from systematically biased beliefs about the effects of medication, to an inadvertent failure to comply caused by forgetfulness or an inability to follow complex treatment plans. Non-adherence can sometimes lead to significant financial costs and physical discomfort, and in some cases, it can cause death, yet these nominally strong incentives are still too weak to motivate “rational” behavior. These reasons for non-adherence closely correspond to the suite of behavioral models that are becoming more common in economics, and the medical literature should make economists more willing to estimate the structural parameters in behavioral models as they seek to scale their results more effectively (DellaVigna, List, and Malmendier 2012; Allcott and Taubinsky 2015).

Second, even as economists acknowledge the apparent prevalence of irrational non-adoption, in their search for countermeasures, they should not expect to find any silver bullets, due to the high levels of context-specificity exhibited by such interventions. This affirms the principle that generalizing results is an imprecise art at best, and that there is no substitute for systematic, incremental research in field contexts that are as close as possible to the target domain. This insight has interesting parallels to the generalizability debate concerning lab experiments (Al-Ubaydli and List 2015a) and calls upon theorists to bridge the important gap between experiment and practice by creating models of generalizability that enhance our fundamental understanding of the where’s, why’s, and how’s of scaling. We trust that this will come down to an understanding of behavioral primitives and how those can be affected as well as learning about features of the environment that attenuate or exacerbate effects of treatment.

This somewhat disheartening result should not be confused with economists having little to learn from the medical literature; learning that interventions that are expected to work are actually unlikely to work constitutes useful information. It potentially saves resources, and allows them to be directed to interventions that are more likely to yield improvements in adherence.

Third, one of the more robust conclusions to emerge from the medical literature is that improvements in adherence are usually the result of complex interventions that combine education, monitoring, and the involvement of other stakeholders. Thus, economists seeking to scale the results of their studies should—after acknowledging the possibility of significant irrational non-adoption—consider skipping simple interventions and going straight to multipronged approaches, especially those that involve exploiting the omnipresence of smartphones, and their interoperability with the electronic systems that underlie the desired modification to behavior.

Moreover, during the evaluation stage, researchers should be careful to focus on the effects of the interventions on the final outcomes associated with the original

proposal for modifying behavior, rather than myopically fixating on the effect of the interventions on adoption, which can sometimes be a red herring. Doctors want patients to get better, not to take pills; equivalently, economists should want people to experience superior outcomes, rather than to modify their behavior as an end in and of itself. In this manner, a bit of backward induction at the design stage can go a long way.

More generally, both clinical researchers and economists stand to gain much from applying experimental design best practices. This includes ensuring that their studies have appropriate sample sizes and sufficient power—see List, Sadoff, and Wagner (2011) for simple rules of thumb for optimal experimental design. In addition, replication and sound inference should be emphasized as countermeasures to the problem of publication bias (Young, Ionnidis, and Al-Ubaydli 2008; Maniadis, Tufano, and List 2014).

Finally, we note one area where economists have shown more initiative than medical researchers: the deployment of financial incentives to improve adoption. There are reasons to be skeptical about the effects of such interventions in the clinical domain, but to the best of our knowledge, the tangible evidence in the medical domain is still limited. Also, one medical study conducted by economists (Charness and Gneezy 2009) is cause for tentative optimism about the beneficial role that financial incentives can play in getting people to overcome the cognitive biases that impede modifications to behavior. Thus, while recommending that economists focus on complex interventions rather than simple ones, we make an exception for the use of financial incentives. These should be systematically compared to nonfinancial alternatives whenever possible for each target context.

■ *We are grateful for comments by Gordon Hanson, Enrico Moretti, and Timothy Taylor.*

References

- Al-Ubaydli, Omar, and John A. List. 2015a. "On the Generalizability of Experimental Results in Economics." Chapter 20 in *Handbook of Experimental Economic Methodology*, edited by G. R. Frechette and A. Schotter. Oxford: Scholarship.
- Al-Ubaydli, Omar, and John A. List. 2015b. "Do Natural Field Experiments Afford Researchers More or Less Control Than Laboratory Experiments?" *American Economic Review* 105(5): 462–66.
- Al-Ubaydli, Omar, John A. List, and Dana L. Suskind. 2017. "What Can We Learn from Experiments? Understanding the Threats to the Scalability of Experimental Results." *American Economic Review* 107(5): 282–86.
- Allais, Maurice. 1990. "Allais Paradox." In *Utility and Probability*, edited by John Eatwell, Murray Milgate, and Peter Newman, 3–9. Springer.
- Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review* 105(8): 2501–38.
- Andreoni, James A., and John H. Miller. 1993.

- "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence." *Economic Journal* 103(418): 570–85.
- Ashraf, Nava, James Berry, and Jesse M. Shapiro.** 2010. "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia." *American Economic Review* 100(5): 2383–2413.
- Benabou, Roland, and Jean Tirole.** 2003. "Intrinsic and Extrinsic Motivation." *Review of Economic Studies* 70(3): 489–520.
- Campbell, Catherine, and Yodwa Mzaidume.** 2002. "How Can HIV Be Prevented in South Africa? A Social Perspective." *British Medical Journal* 324(7331): 229–32.
- Caplan, Bryan.** 2002. "Systematically Biased Beliefs about Economics: Robust Evidence of Judgmental Anomalies from the Survey of Americans and Economists on the Economy." *Economic Journal* 112(479): 433–58.
- Carroll, Christopher D., Jody Overland, and David N. Weil.** 2000. "Saving and Growth with Habit Formation." *American Economic Review* 90(3): 341–55.
- Chang, Eric C., Joseph W. Cheng, and Ajay Khorana.** 2000. "An Examination of Herd Behavior in Equity Markets: An International Perspective." *Journal of Banking and Finance* 24(10): 1651–79.
- Charness, Gary, and Uri Gneezy.** 2009. "Incentives to Exercise." *Econometrica* 77(3): 909–31.
- Chopra, Mickey, and Neil Ford.** 2005. "Scaling up Health Promotion Interventions in the Era of HIV/AIDS: Challenges for a Rights Based Approach." *Health Promotion International* 20(4): 383–90.
- Dayer, Lindsey, Seth Heldenbrand, Paul Anderson, Paul O. Gubbins, and Bradley C. Martin.** 2013. "Smartphone Medication Adherence Apps: Potential Benefits to Patients and Providers." *Journal of the American Pharmacists Association*: 53(2): 172–81.
- Deaton, Angus, and Nancy Cartwright.** 2016. "Understanding and Misunderstanding Randomized Controlled Trials." NBER Working Paper 22595.
- Deci, E. L., R. Koestner, and R. M. Ryan.** 1999. "A Meta-analytic Review of Experiments Examining the Effects of Extrinsic Rewards on Intrinsic Motivation." *Psychology Bulletin* 125(6): 627–68.
- DellaVigna, Stefano, John A. List, and Ulrike Malmendier.** 2012. "Testing for Altruism and Social Pressure in Charitable Giving." *Quarterly Journal of Economics* 127(1): 1–56.
- Dennison, Laura, Leanne Morrison, Gemma Conway, and Lucy Yardley.** 2013. "Opportunities and Challenges for Smartphone Applications in Supporting Health Behavior Change: Qualitative Study." *Journal of Medical Internet Research* 15(4).
- Duflo, Esther, Michael Kremer, and Jonathan Robinson.** 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101(6): 2350–90.
- Farnan, J. M., L. A. Petty, E. Georgitis, S. Martin, E. Chiu, M. Prochaska, and V. M. Arora.** 2012. "A Systematic Review: The Effect of Clinical Supervision on Patient and Residency Education Outcomes." *Academic Medicine* 87(4): 428–42.
- Ferraro, Paul J., and Michael K. Price.** 2013. "Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment." *Review of Economics and Statistics* 95(1): 64–73.
- Field, Erica, and Rohini Pande.** 2008. "Repayment Frequency and Default in Microfinance: Evidence from India." *Journal of the European Economic Association* 6(2–3): 501–09.
- Finitsis, David J., Jennifer A. Pellowski, and Blair T. Johnson.** 2014. "Text Message Intervention Designs to Promote Adherence to Antiretroviral Therapy (ART): A Meta-analysis of Randomized Controlled Trials." *PLOS One* 9(2): e88166.
- Fryer, Roland G., Jr.** 2011. "Financial Incentives and Student Achievement: Evidence from Randomized Trials." *Quarterly Journal of Economics* 126(4): 1755–98.
- Fryer, Roland G., Jr., Steven D. Levitt, and John A. List.** 2015. "Parental Incentives and Early Childhood Achievement: A Field Experiment in Chicago Heights." NBER Working Paper 21477.
- Garrett, Laurie.** 2007. "The Challenge of Global Health." *Foreign Affairs* 86(1): 14–38.
- Gill, Preetinder S., Ashwini Kamath, and Tejkar S. Gill.** 2012. "Distraction: An Assessment of Smartphone Usage in Health Care Work Settings." *Risk Management and Healthcare Policy* 5: 105–14.
- Giuffrida, Antonio, and David J. Torgerson.** 1997. "Should We Pay the Patient? Review of Financial Incentives to Enhance Patient Compliance." *BMJ* 315(7110): 703–07.
- Gosnell, Greer K., John A. List, and Robert Metcalfe.** 2016. "A New Approach to an Age-Old Problem: Solving Externalities by Incenting Workers Directly." NBER Working Paper 22316.
- Halpern, S. D., Benjamin French, Dylan Small, Kathryn Saulsgiver, Michael O. Harhay, Janet Audrain-McGovern, George Loewenstein, Troyen A. Brennan, David A. Asch, and Kevin G. Volpp.** 2015. "Randomized Trial of Four Financial-Incentive Programs for Smoking Cessation." *New England Journal of Medicine* 372(22): 2108–17.
- Haynes, R. B., E. Ackloo, N. Sahota, H. P. McDonald, and X. Yao.** 2008. "Interventions for Enhancing Medication Adherence." *Cochrane*

Database Systematic Reviews 16(2).

Hibbard, Judith H., Jean Stockard, and Martin Tusler. 2005. "Hospital Performance Reports: Impact on Quality, Market Share, and Reputation." *Health Affairs* 24(4): 1150–60.

Hossain, Tanjim, and John A. List. 2012. "The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations." *Management Science* 58(12): 2151–67.

Huang, H. L., Y. C. Li, Y. C. Chou, Y. W. Hsieh, F. Kuo, W. C. Tsai, S. D. Chai, B. Y. Lin, P. T. Kung, and C. J. Chuang. 2013. "Effects of and Satisfaction with Short Message Service Reminders for Patient Medication Adherence: A Randomized Controlled Study." *BMC Medical Informatics and Decision Making* 13(1): 127.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies." *Journal of Economic Perspectives* 5(1): 193–206.

Klemperer, Paul. 1987. "Markets with Consumer Switching Costs." *Quarterly Journal of Economics* 102(2): 375–94.

Kreuter, Mathew, David Farrell, Laura Olevitch, and Laura Brennan. 2000. *Tailoring Health Messages: Customizing Communication with Computer Technology*. New York: Routledge.

Kripalani, S., X. Yao, and R. B. Haynes. 2007. "Interventions to Enhance Medication Adherence in Chronic Medical Conditions: A Systematic Review." *Archives of Internal Medicine* 167(6): 540–50.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2): 443–77.

Lavorini, F., A. Magnan, J. C. Dubus, T. Voshaar, L. Corbetta, M. Broeders, R. Dekhuijzen, J. Sanchis, J. L. Viejo, P. Barnes, C. Corrigan, M. Levy, and G. K. Crompton. 2008. "Effect of Incorrect Use of Dry Powder Inhalers on Management Patients with Asthma and COPD." *Respiratory Medicine* 102(4): 593–604.

Levitt, Steven D., and John A. List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *Journal of Economic Perspectives* 21(2): 153–74.

List, John A., Sally Sadoff, and Mathis Wagner. 2011. "So You Want to Run an Experiment, Now What? Some Simple Rules of Thumb for Optimal Experimental Design." *Experimental Economics* 14(4): 439–57.

Lu, C., C. M. Michaud, K. Khan, and C. J. Murray. 2006. "Absorptive Capacity and Disbursements by the Global Fund to Fight AIDS, Tuberculosis and Malaria: Analysis of Grant Implementation." *Lancet* 368(9534): 483–88.

Maniadis, Zacharias, Fabio Tufano, and John A.

List. 2014. "One Swallow Doesn't Make a Summer: New Evidence on Anchoring Effects." *American Economic Review* 104(1): 277–90.

McDonald, H. P., A. X. Garg, and R. B. Haynes. 2002. "Interventions to Enhance Patient Adherence to Medication Prescriptions: Scientific Review." *Journal of the American Medical Association* 288(22): 2868–79.

Nieuwlaat, R., N. Wilczynski, T. Navarro, N. Hobson, R. Jeffery, A. Keepanasseril, T. Agoritsas, N. Mistry, A. Iorio, S. Jack, B. Sivaramalingam, E. Iserman, R. A. Mustafa, D. Jedraszewski, C. Cotoi, and R. B. Haynes. 2014. "Interventions for Enhancing Medication Adherence." *Cochrane Database Systematic Review* 20(11).

Orr, R. D., N. Pang, E. D. Pellegrino, and M. Siegler. 1997. "Use of the Hippocratic Oath: A Review of Twentieth Century Practice and a Content Analysis of Oaths Administered in Medical Schools in the U.S. and Canada in 1993." *Journal of Clinical Ethics* 8(4): 377–88.

Paulsell, Diane, Sarah Avellar, Emily Sama Martin, and Patricia Del Grosso. 2010. *Home Visiting Evidence of Effectiveness Review: Executive Summary*. Princeton, NJ: Mathematic Policy Research.

Peterson, A. M., L. Takiya, and R. Finley. 2003. "Meta-analysis of Trials of Interventions to Improve Medication Adherence." *American Journal of Health-System Pharmacy* 60(7): 657–65.

Petrie, K. J., K. Perry, E. Broadbent, and J. Weinman. 2012. "A Text Message Programme Designed to Modify Patients' Illness and Treatment Beliefs Improves Self-Reported Adherence to Asthma Preventer Medication." *Journal of Health Psychology* 17(1): 74–84.

Quinn, C. C., S. S. Clough, J. M. Minor, D. Lender, C. C. Okafor, and A. Gruber-Baldini. 2008. "WellDoc Mobile Diabetes Management Randomized Controlled Trial: Change in Clinical and Behavioral Outcomes and Patient and Physician Satisfaction." *Diabetes Technology and Therapeutics* 10(3): 160–68.

Raikes, Helen, Beth L. Green, Jane Atwater, Ellen Kisker, and Jill Constantine. 2006. "Involvement in Early Head Start Home Visiting Services: Demographic Predictors and Relations to Child and Parent Outcomes." *Early Childhood Research Quarterly* 21(1): 2–24.

Roggman, Lori A., Gina A. Cook, Carla A. Peterson, and Helen H. Raikes. 2008. "Who Drops out of Early Head Start Home Visiting Programs?" *Early Education and Development* 19(4): 574–99.

Sarabi, Roghayeh Ershad, Farahnaz Sadoughi, Roohangiz Jamshidi Orak, and Kambiz Bahaadinbeigy. 2016. "The Effectiveness of Mobile Phone Text Messaging in Improving Medication Adherence for Patients with Chronic Diseases: A

Systematic Review." *Iranian Red Crescent Medical Journal* 18(5).

Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment." *Review of Economic Studies* 71(2): 513–34.

Simon, H. A. 1972. "Theories of Bounded Rationality." In *Decision and Organization*, edited by C. B. McGuire and Roy Radner, 161–76. Amsterdam: North Holland Publishing Company.

Smith, Aaron. 2015. *U.S. Smartphone Use in 2015*. Washington, DC: Pew Research Center.

Strandbygaard, U., S. F. Thomsen, and V. Backer. 2010. "A Daily SMS Reminder Increases Adherence to Asthma Treatment: A Three-Month Follow-up Study." *Respiratory Medicine* 104(2): 166–71.

Studdert, David M., Michelle M. Mello, Atul A. Gawande, Tejal K. Gandhi, Allen Kachalia, Catherine Yoon, Ann Louise Puopolo, and Troyen A. Brennan. 2006. "Claims, Errors, and Compensation Payments in Medical Malpractice Litigation." *New England Journal of Medicine* 354(19): 2024–33.

Telyukova, Irina A., and Randall Wright. 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle." *Review of Economic Studies* 75(2): 629–47.

Thaler, Richard H., and Shlomo Benartzi. 2004. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving." *Journal*

of Political Economy 112(1): S164–87.

Venkatraman, A., N. Garg, and N. Kumar. 2015. "Greater Freedom of Speech on Web 2.0 Correlates with Dominance of Views Linking Vaccines to Autism." *Vaccine* 33(12): 1422–25.

Vervloet, M., A. J. Linn, J. C. van Weert, D. H. de Bakker, M. L. Bouvy, and L. van Dijk. 2012. "The Effectiveness of Interventions Using Electronic Reminders to Improve Adherence to Chronic Medication: A Systematic Review of the Literature." *Journal of American Medical Informatics Association* 19(5): 696–704.

Volpp, K. G., L. K. John, A. B. Troxel, L. Norton, J. Fassbender, and G. Loewenstein. 2008. "Financial Incentive-Based Approaches for Weight Loss: A Randomized Trial." *Journal of the American Medical Association* 300(22): 2631–37.

Watt, Richard, Francisco J. Vazquez, and Ignacio Moreno. 2001. "An Experiment on Rational Insurance Decisions." *Theory and Decision* 51(2–4): 247–96.

Young, Neal S., John P. A. Ioannidis, and Omar Al-Ubaydli. 2008. "Why Current Publication Practices May Distort Science." *PLOS Medicine* 5(10): e201.

Zullig, L. L., E. D. Peterson, and H. B. Bosworth. 2013. "Ingredients of Successful Interventions to Improve Medication Adherence." *Journal of the American Medical Association* 309(24): 2611–12.

