NBER WORKING PAPER SERIES

NON-ADHERENCE IN HEALTH CARE: A POSITIVE AND NORMATIVE ANALYSIS

Mark Egan
Tomas J. Philipson

Working Paper 20330
http://www.nber.org/papers/w20330

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2014

We are thankful to seminar participants at several institutions as well as Michael Dickstein, Dana Goldman, Kris Hult, Alex MacKay and Derek Neal for discussions and comments. Tomas Philipson is thankful for financial support from the George Stigler Center for Study of the Economy and the State, The University of Chicago. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Mark Egan and Tomas J. Philipson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Non-Adherence In Health Care: A Positive and Normative Analysis
Mark Egan and Tomas J. Philipson
NBER Working Paper No. 20330
July 2014
JEL No. I1,I10,I18

ABSTRACT

Non-adherence in health care results when a patient does not initiate or continue care that has been recommended by a provider. Previous researchers have identified non-adherence as a major source of waste in US healthcare, totaling approximately 2.3% of GDP, and have proposed a plethora of interventions to improve adherence. However, little explicit analysis exists in health economics of the dynamic demand behavior that drives non-adherence. We argue that while providers may be more informed about the population-wide effects of treatments, patients are more informed about their individual treatment effect. We interpret a patient's adherence decision as an optimal stopping problem where patients learn the value of a treatment through experience. Our positive analysis derives an "adherence survival function" and shows how various observable factors affect adherence. Our normative analysis derives the efficiency effects of non-adherence, the conditions under which adherence is too high or too low, and why many common interventions aimed at raising adherence produce indeterminate welfare effects. We calibrate these welfare effects for one of the largest US drug categories, cholesterol reducing drugs. Contrary to frequent normative claims of under-adherence, our estimates suggest that the ex-post efficiency loss from over-adherence is over 80% larger than from under-adherence.

Mark Egan
University of Chicago
Department of Economics
1126 East 59th Street
Chicago, IL 60637
egan@uchicago.edu

Tomas J. Philipson
Irving B. Harris Graduate School
of Public Policy Studies
University of Chicago
1155 E. 60th Street
Chicago, IL 60637
and NBER
t-philipson@uchicago.edu

Section 1: Introduction

A wide-spread challenge in health care is getting patients to adhere to prescribed treatments and therapies. In the United States, estimates show that non-adherence is wasteful[2]; the New England Healthcare Institute (2009) estimate that the annual cost of non-adherence in the U.S is approximately $290 billion, equating to about 13% of total health care spending or 2.3% of GDP. Therefore, improving medical adherence through both private and public interventions has been identified as a crucial step to improving health outcomes and lowering health care costs.[3] Recent technological advancements have targeted medical adherence such as electronic and educational messaging systems (Baum 2013, Comstock 2013, Vollmer et al. 2011), as well as technology designed to help providers to identify non-adherent patients (Lesselroth et al. 2011). In the US, The Patient-Centered Outcomes Research Institute (PCORI), the Agency for Healthcare Research and Quality, and the National Institutes of Health among other government bodies, have all dedicated substantial funding to support research on raising medication adherence (Pharmaceutical Research and Manufacturers of America, 2013). These efforts have been driven by an enormous literature outside of economics on the prevalence of non-adherence and its consequences. Indeed, since 1996 it is estimated that more than 25,000 peer-reviewed medical articles have been published on patient-adherence or compliance (Chernew 2008). The overall implicit concern of this vast literature is that adherence is too low and that private and public interventions are needed to raise adherence.

---

[2] See Bosworth et al. 2011 for further discussion.
[3] See Black et al. (1987), Feldman et al. (1998), Flack et al. (1996), Haynes et al. (1996), Hershey et al. (1980), Mallion et al. (1998), and Nelson et al. (1980).

However, there exists little explicit economic analysis of the dynamic demand behavior resulting in non-adherence that offers predictions about the conditions under which it is more likely to occur than not. Without an empirical validation of such a positive theory, it is of course difficult to make credible normative claims about whether adherence is too low or too high. To meet this end, this paper provides an explicit analysis of non-adherence and derives its positive and normative implications.

We interpret non-adherence as an optimal stopping problem for a patient learning about his individual value of the therapy. Although providers recommending treatments are likely more informed about the population-wide effects of these treatments, patients experiencing a treatment are more informed about its individual specific value. This individual value of treatment incorporates how the patient trades off patient specific treatment effectiveness, side effects, and costs of care. In our analysis a patient's prior beliefs about a treatment coupled with the patient's experience with the treatment drive initiation and subsequent adherence. The patient behavior mimics the common sense practice of going on a treatment, seeing if it is valued by them, and terminating it if it is not. Indeed, non-adherence is an inherently a dynamic demand behavior that requires an explanation of why people initiate but then discontinue therapy. Our positive analysis of non-adherence as an optimal stopping problem offers many testable implications. As patients learn about the treatment, they will eventually become more informed over time. This implies that good matches of patients to treatments last but bad ones do not. More precisely, we derive an "adherence survival function" depicting the share of patients still on treatment as a function of time, and show how various observable factors affect adherence. We predict that non-adherence occurs early in the sense that adherence decisions stabilize with sufficient learning about treatment value. We also predict that education has non-trivial effects on adherence because it

interacts with patient level treatment effects; more educated individuals adhere longer to valuable care, but shorter to what turns out to be invaluable care for them.[4] In addition, we predict that the quality of providers and their communication with patients are likely to impact short-run rather than long-run adherence behavior.

Non-adherence driven by patient-learning has strong normative implications. Although there are many analysts stressing inadequate adherence in both economic and medical circles, claims of under-consumption made by bystanders are traditionally viewed with skepticism. In particular, we argue it is important to separate between ex-ante versus ex-post efficient adherence. When learning about personalized treatment value, patients act in an ex-ante optimal fashion given their treatment beliefs. However, adherence may be ex-post inefficient in that some patients do not adhere to what turns out to be a valuable therapy for them, while others adhere to what turns out to be a non-valuable therapy. Therefore there is over-adherence by those who do not respond to therapy and under-adherence by those who respond. We argue that over-adherence vanishes over time as patients eventually learn that they do not respond to the therapy. However, under-adherence is permanent because we show that once a patient leaves the treatment he will never find it optimal to re-adhere. We argue this asymmetry of welfare losses makes it non-trivial to evaluate the efficiency effects of many interventions aimed at raising adherence, as they customarily raise adherence of both responders and non-responders.

We calibrate these efficiency effects in the case of the cholesterol reducing drug simvastatin (Zocor). Interestingly, our calibration results imply that the vast majority of the efficiency loss comes from over-adherence, as opposed to under-adherence, even though less

[4] Such effects are the analog to uneducated individuals adhering more to smoking after it was discovered that cigarettes had the "side-effect" of inducing cancer.

than half of patients adhere. Specifically, we find that the ex-post efficiency loss from over-adherence is over 80% larger than that from under-adherence. In this context, we stress that simply arguing that there is too little adherence to cholesterol reducing therapies because patients do not understand the treatment benefits seems unsatisfactory as those perceived benefits presumably made them start the therapy in the first place.

This paper relates to several strands of previous analysis. There is of course a large literature on health care demand including that of Grossman (1972), but we are not aware of any explicit analysis of the dynamic demand behavior that is inherent in non-adherence. Elsewhere (Seabury et al. 2014), we have provided a partial review of the vast empirical health services research literature on the extent of non-adherence. This paper may be viewed as the direct post-approval analog of Philipson and Hedges (1998) and Philipson and Desimone (1997) who analyze the effects of attrition in clinical trials when subjects learn about individual treatment effects in a similar manner to how investigators learn about population-wide effects. The paper also relates to the structural estimation of Dickstein (2014) but differs in drawing out the positive and normative implications of the optimal stopping problem inherent in non-adherence. Goldman et al (2007) and Chernew et al (2008) report negative price elasticities for this type of demand. In the general economics literature, this paper is most closely related in spirit and structure to the labor literature on job-turnover (Jovanovic 1979) where matching workers to jobs is the analog of matching patients to treatments.

The paper is briefly outlined as follows. Section 2 discusses a simple illustrative case of non-adherence. Section 3 provides the general properties of the adherence survival function. Section 4 discusses the large set of positive implications regarding the effects of observable factors on the adherence survival function. Section 5 discusses the normative implications for

efficient adherence. Section 6 calibrates the size of efficiency effects for the cholesterol reducing drug simvastatin. Section 7 generalizes our analysis to multiple treatments. Lastly, Section 8 concludes with a discussion of several future research avenues suggested by the explicit analysis of this type of demand behavior.

Section 2: A Simple Illustration of Non-Adherence as an Optimal Stopping Problem

To illustrate some of the ideas in the simplest possible fashion, we first discuss a two period example where a patient is learning whether or not he responds to a treatment. The patient decides whether to initiate the treatment regime with limited information regarding the value of the treatment. By initiating treatment, he learns whether he responds to the treatment and then decides to continue to adhere or stop the treatment. We are interested in the observable conditions under which so called *primary* adherence occurs, the patient initiates the treatment, as well as when *secondary* adherence occurs, the patient continues after initiation.

Consider a treatment regime where the patient either responds or not to treatment giving rise to the health outcomes levels denoted $h_1 \geq h_0$. The health outcome throughout the paper is interpreted as a *net benefit* index of all relevant health aspects of a treatment, inclusive of treatment effectiveness, side effects, and any other effects on a patient's health. The price of the treatment is denoted p, which would be the full cost including co-pays, time costs of compliance or travel, and any other cost factors. Treatment yields utility $U_1 = U(h_1, p)$ if the patient is a responder and utility $U_0 = (h_0, p)$ if the patient is a non-responder. Utility is distinct from the

effectiveness or health; effective treatments may have little value and low adherence by a patient not concerned with the condition being treated.

Prior to treatment, patients do not know if they are responders or non-responders. Patients have the outside option of forgoing the treatment and undertake the standard of care which yields utility $U_s = U(h_s, p_s)$. We assume that $U_1 > U_s > U_0$ so that responders would like to undertake the treatment while non-responders will prefer the standard of care.

To understand the primary adherence decision we must first understand the secondary adherence decision. After the patient has learned whether they value the therapy responders adhere if

$$U_1 \geq U_s$$

The patient stops therapy (non-adheres) if he learns he does not value it

$$U_0 < U_s$$

Let F denote the patient's prior belief that he values the treatment which we treat as being a responder. The prior belief F may or may not correspond to the true proportion responding in the population, denoted F^*. The patient optimally initiates treatment whenever

$$[FU_1 + (1 - F)U_0 - U_s] + \beta\,[F(U_1 - U_s)] \geq 0$$

The first term has an indeterminate sign and is the difference between the uncertain value of the new therapy and the standard of care during the first period. The second term is positive and represents the optional value from learning that the treatment is valuable and therefore adhering to it in the future. The second term is discounted by the factor β. If the first term is positive, primary adherence occurs as there is no downside to it. But even if the patient does not initially believe in the treatment he may adhere to it due to the option value of learning the treatment is indeed valuable. Overall, the health impact of the treatments involved and the utilities and beliefs

of patients, represented by the distribution of (U_s, U_1, U_0, F) in the patient population, determine the rates of primary and secondary adherence in the population.

This simple case reveals some basic implications that are useful to illustrate the general analysis later. First, clearly a larger price or copay for the treatment, interpreted as a larger p, lowers both primary and secondary adherence. Demand is downward sloping.

Second, the impact of competition in the treatment class may be represented by a higher quality or lower price of the standard of care, both raising the outside option U_s. This both raises the cost of learning through primary adherence (first term) and lowers the benefit of learning (second term). Consequently, competitive classes are predicted to have lower treatment specific adherence.

Third, a provider may have expertise in the *population-wide effects* of the treatment, here the fraction of responders F^*, while the patient has expertise in the *individual value* of the treatment after experiencing care. The individual value is partly but not fully determined by the personalized treatment effect. Providers will not have expertise in how a given patient trades off various aspects of the treatment after learning about it, such as side-effects, efficacy or price. In particular, doctors may argue that patients under-adhere because doctors are often only focused on health outcomes as opposed to the patients that must weigh all aspects of care and pay the price p.

Fourth, communication between providers and patients may greatly affect adherence. The doctor may through communication influence the prior of the patient to reduce a potential misperception of the population wide effect, F-F^*, but ultimately treatment experience outweighs prior beliefs in driving adherence. Nevertheless, there is a predicted link between primary and secondary adherence driven by the priors but they affect primary adherence more

than secondary adherence. Even without communication, providers may greatly affect adherence. For example, higher quality doctors may have patients with higher adherence because the patient trusts the doctor's opinion more which thereby raises the patient's belief in the treatment recommended by the provider. The patient may even act against his own initial belief F if learning that the doctor recommends the treatment, thereby updating beliefs closer to F^*. However, in the simplest case here, the patient learns the value of the treatment directly, which means that only primary and not secondary adherence is influenced by the provider. This generalizes later on, as it will be the case that the patient ultimately learns from his experience whether a treatment is valuable. Thus, providers may drive short-run or primary adherence but may have limited impact on long-run secondary adherence when the experience of being on treatment matters more. In the extreme case above the true value of the treatment is learned by the patient immediately, so priors have no effect on secondary adherence.

Fifth, the true treatment effect, that is, the fraction F^* of patients who actually respond to the treatment, affects adherence whether or not this true effect is close or not to the prior F of that treatment effect. The true treatment effect raises secondary adherence because the fraction of responders in the population increases. Put simply, if non-adherence is due to non-response, more favorable treatment effects raise adherence. Therefore, ceteris paribus, the adherence behavior of the same individual may differ across treatments according to how that individual values the different treatments and the performance he experiences on them. For individual level data on health outcomes and adherence, differences in the observed and unobserved outcomes may drive differential adherence across treatments. For example, a patient may adhere perfectly to a pain medicine while adhering poorly to a cholesterol reducing drug given the immediate benefit of the former and delayed benefit of the latter. Overall, population-wide effectiveness

will be positively related to population level adherence as positive treatment experience leads to continued adherence.

Sixth, education interacts with individual treatment effects in driving adherence. One way of interpreting treatment education is having a prior belief F closer to the true proportion responders F^*. Under this interpretation, education may increase or lower adherence since beliefs rather than accuracy of those beliefs drives adherence. For example, if poorly educated individuals are too optimistic about natural medicines with no benefits, they may adhere more than educated individuals. On the other hand, if educated individuals understand the benefits of adherence to traditional medicines better than uneducated individuals they may adhere more.

Seventh, later on we will discuss the "survival function" of adherence that measures the fraction still adhering as a function of the length of treatment. More precisely, this survival results from the result that here is no "re-adherence" after the patient learns the treatment is not valuable. In the extreme case of immediate learning discussed above, the survival function is a step function with the step-size being driven by the true non-response rate $1 - F^*$. This illustrates the more general theme that non-adherence tends to occur early on and then slow down. The hazard rate of non-adherence goes to zero with sufficient learning about treatment value. Once patients learn the true value of treatment, only true responders will adhere.

Eighth, the so called "cost-effectiveness" of care is loosely related to larger adherence. Here, cost-effectiveness is $[Fh_1 + (1 - F)h_0]/p$ and we already discussed why a rise in price lowers adherence and a rise in effectiveness generally raises it. Cost-effectiveness is not perfectly related to adherence for several reasons. One reason is that patients may trade off cost versus effectiveness differently than one to one, that is, the utility $U(h, p)$ may differ from the ratio.

Other reasons include unmeasured quality dimensions that affect patient utility and adherence, or prices that do not correspond to the full cost of care faced by the patient.

Ninth, treatment duration effects on adherence may be non-standard. Consider when the single period of learning in this illustrative case may involve different amounts of calendar time. For example, it may take many weeks before any learning takes place through lab tests indicating treatment quality, as e.g. would be the case for statin therapy affecting cholesterol levels. Consequently learning takes place at a slower rate. Similarly, a patient may be more willing to try out a therapy for a chronic versus temporary condition because of the larger option value of the former.

Tenth, a patient's willingness to trade-off current versus future health benefits affects adherence. The discount factor (β) lowers the option value of adherence. Therefore, adherence will be higher for treatments with immediate benefits, e.g. pain or fever reduction, than delayed benefits, e.g. lowered cholesterol levels that affect future mortality. In addition, other behavior that also is driven by discounting, such as e.g. savings behavior, education, or other forms of health investments, may be positively correlated with adherence while not causing it directly.

Eleventh, patient misinformation or bias, in the sense of $F - F^*$ being non-zero making the patient too optimistic or pessimistic about the treatment, is neither necessary nor sufficient for non-adherence. In fact, in the example above, regardless of what the patient thought about the treatment before he started, as long as he was optimistic enough to initiate the care, he would not adhere after learning that the treatment was not valuable to him. Biases about treatment effects must, in addition, be weighed against price.

Lastly, non-adherence through learning has strong normative implications. The patient's in the example above cannot be made better off in an ex-ante sense, since they adhere optimally

given their beliefs. However, there may be some patients that may not engage in primary adherence even though they would have benefited from treatment and there may be some patients that adhere even though they ultimately do not benefit from treatment. Note that in the case above if initially everyone adheres but then non-responders quit, the entire ex-post efficiency loss is due to over-adherence, not under-adherence as often argued. In later sections we derive more precisely the properties of such ex-post inefficiencies resulting from non-adherence induced by learning.

Section 3: Non-Adherence as an Optimal Stopping Problem

We here extend the simple illustrative case to multiple periods and a larger set of individual treatment responses. We derive the implications of interpreting non-adherence as an optimal stopping problem when learning about one's own personal treatment effect given prior knowledge about population wide effects.[5]

We now assume that there is a continuum of patient types or true treatment effects denoted by q and distributed according to $F(\cdot)$. These treatment effects correspond to the patient-specific "quality" of the product that that the treatment represents. The health of the patient in a given period reflects the quality of the treatment plus some idiosyncratic shock (noise) according to

$$h_t = q + \varepsilon_t$$

The observable health of the patient depends on the unobservable quality of treatment q_t as well as other factors, ε_t. The health of the patient may be determined by other factors than the quality

[5] Our model is similar in spirit to Jovanovic's (1979) and Ljungqvist and Sargent (2006) analysis of job turnover.

of treatment, such as the ability of the body itself to cure a disease. Thus, the patient cannot infer treatment quality immediately but learns over time. For example, pain or fever reduction may be due to a treatment or the body healing naturally.

The period utility benefit for treatment is given by $U(h_t, p)$ where

$$U(h_t, p) = h_t - \gamma p$$

The parameter γ represents patient's health consumption trade-off. As before we assume that patients have access to an alternative standard of care with per period utility $U(h_s, p_s)$.

Patients have a prior over the quality of the treatment, $F_0(\cdot)$. For now we assume that each patient's initial prior beliefs reflect the true distribution of treatment heterogeneity such that $F_0(\cdot) = F(\cdot)$. This may be interpreted as patients agreeing with providers about the population wide effects of treatments before learning about their individual value of care. For example, this prior may be the result of knowing summary statistics of the distribution of treatment effects from the labeling of the product, obtained from clinical trials in the approval process. We denote a patient's prior at time t as $F_t(\cdot|\vec{h}_t)$ where $\vec{h}_t$ is the history of personal health outcomes on the treatment. Under the maintained assumption that the prior and shocks are normally distributed, $F_0(q) \sim \ N(\mu_0, \sigma_0)$ and noise is distributed $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, standard normality results imply that a patient's posterior distribution over their treatment effect is given by

$$F_t(q_i|\vec{h}_t) = N\left(\omega_t \bar{h}_t + (1-\omega_t)\mu_0, \left[\frac{t}{\sigma_\varepsilon^2} + \frac{1}{\sigma_0^2}\right]^{-1}\right), \ \omega_t = \frac{\sigma_0^2}{\frac{\sigma_\varepsilon^2}{t} + \sigma_0^2} \tag{1}$$

The patient optimally updates his beliefs about the quality of the treatment based on weighting the average health outcomes, $\bar{h}_t$, and his initial prior. With each observation the patient places more weight on his treatment experience and less on his prior. In addition, his posterior variance

decreases after each treatment experience over time. In other words, the longer a patient has been in treatment, the more he learns about the quality of the treatment and the more his belief is informed by his own experience rather than any beliefs prior to initiating the treatment, such as the population-wide beliefs offered by the provider.

Given these beliefs about the personal treatment effect, the value function after t rounds of treatments is given by

$$V(h_t, \vec{h}_t, F_0) = U(h_t, p) + \beta \max \left\{ E[V(h_{t+1}, \vec{h}_{t+1}, F_0)|F_t], \frac{U(h_s, p_s)}{1-\beta} \right\}$$

The patient elects to adhere to treatment only if the expected value of staying on the treatment is larger than stopping and going on the standard of care from there on. Once a patient elects to forgo treatment for the standard of care, he will find it optimal to continue the standard care in all proceeding periods.[6] Patients will never find it optimal to "re-adhere" to the treatment regime. As before, because of the future option value of continuing treatment, a consumer may elect adhere to treatment even if the expected future period return is lower than that of the standard of care.

It is well established that the optimal stopping behavior for this type of learning is characterized by treatment performance threshold (Gittins and Jones 1974, Gittins and Jones 1979).[7] This implies that non-adherence occurs when the average experience on treatment $\bar{h}_t$ is below a certain threshold level, here denoted z_t. Adherence behavior conditional on patient type

---

[6] Suppose patients potentially found it optimal to reenter treatment. Consider the patients decision to continue treatment after receiving n rounds of treatment

$$max \left\{ E[V(h_{n+1}\vec{h}_{n+1}, F_0)|F_n], V(h_s, \vec{h}_n, F_0) \right\}$$

If a patient opts for the standard of care he does not learn any additional information about the treatment regime. Thus if a patient opts for the standard of care he will do so in all proceeding periods. This was originally shown in Bradt, Johnson and Karlin (1956).

[7] See also Gittins et. al (2011) and Powell and Ryzhov (2012) for a general discussion of characterizing stopping problems.

(q) is characterized by survival function $S(t|q)$ which reflects the proportion of type q individuals remaining in treatment at time t

$$S(t|q) = \Pr(h_1 > z_1, \bar{h}_2 > z_2, ..\bar{h}_t > z_t|q)$$

The overall survival function of adherence results from aggregating over all types

$$S(t) = \int S(t|q)dF(q)$$

For such a survival function, the degree of primary non-adherence or non-initiation corresponds to the magnitude $1 - S(1)$ while secondary non-adherence or discontinuation corresponds to $S(1) - S(t)$. In the simple illustrative case discussed in the previous section, assuming everyone initially adheres the survival function was a step-function with step size equaling the share of non-responders; $S(1) = 1, S(2) = F^*$.

Section 4: Positive Implications about Factors Driving Non-Adherence

In this section we discuss the many testable implications of interpreting non-adherence as optimal stopping when patients learn about individual treatment value.

4.1 Cost of Care and Adherence

There are two primary costs of treatment: the monetary cost of treatment p and the opportunity cost of treatment. The opportunity cost of treatment is the cost incurred by forgoing the existing standard of care. The cost of treatment reduces the net benefit of treatment and ultimately lowers adherence. In the context of the preceding model, the threshold average experience required for patients remain adherent at time t, z_t, is increasing in the cost of treatment; $\frac{dz_t}{dp} > 0$. This rise in price may be either due to higher co-pays, premiums, or other

forms of time or monetary costs that contribute to the total cost of care. For example, adherence should rise with patent expirations if they lower prices.

The opportunity cost in terms of the standard of care, or other substitute forms of care, represents the outside option in our stopping problem. Thus, it is straightforward to show that the better the outside options, the lower the adherence. This implies that the price of the standard of care raises adherence while the quality of standard of care lowers it; $\frac{dz_t}{dp_s} < 0$, $\frac{dz_t}{dq_s} > 0$. Thus, external market forces outside the given treatment help determine adherence. For example, generic or therapeutic competition for the standard of care lowers adherence for alternative treatments, and the relative price subsidies across treatments in a class affect adherence.

4.2 Treatment Quality and Adherence

A basic implication of our analysis is that better treatment experiences leads to higher adherence both on the individual personalized level as well as on an aggregate level relating to overall product quality.

4.2.1 Individual Treatment Performance and Adherence

A fundamental implication of our analysis is that, other things constant, patients who perform better on a treatment adhere longer to it. In other words, good matches of patients to treatments last and bad matches do not. In our framework, the most basic way in which this occurs is that if a patient experiences treatment outcomes $\vec{h}$ that are uniformly larger than another set of treatment outcomes $\vec{h'}$, then he will adhere longer to the first. In our particular learning environment based on normality assumptions, the first set of experiences would imply a larger average health outcome throughout, which in turn would imply a higher posterior mean, thus resulting in higher adherence.

4.2.2 Aggregate Treatment Performance

Differences across treatments in terms of their overall quality, such as e.g. effectiveness and lack of side effects, are represented by differences in the mean quality of the treatment. These population-wide effects of treatments are often estimated in clinical trials conducted to gain approval for marketing. At any given time, the threshold driving the optimal stopping rule is decreasing in the prior mean/average effectiveness of the treatment regime.[8]

$$\frac{dz_t}{d\mu_0} \leq 0$$

In fact, an increase in the average effectiveness of the treatment raises adherence through two channels. It lowers the threshold above which adherence occurs but in addition it increases the probability the average health outcome exceeds the stopping rule threshold. If the impact of the treatment is measured in, for example, how significantly it raises quality-adjusted-life-years (QALY), then this prediction says that ceteris paribus patients are more adherent to treatments the more they raise those measured QALY levels. However, if higher quality treatments also are more expensive, then the quality and price effects counteract each other.

4.3: Treatment Heterogeneity and Noise

A patient's health outcome in a given period is a function of the unknown treatment quality ($q \sim F(\mu, \sigma^2)$) and the treatment/signal noise ($\varepsilon_t \sim F(0, \sigma_\varepsilon^2)$). Treatment heterogeneity, captured by the prior variance σ^2, and the signal noise, captured by σ_ε^2 , impact the value of learning. Treatment heterogeneity increases the option value of learning while signal noise

[8]See section 6.4 in Gittins (1989) and Corollary 1 in Yao (2006) for further details.

constrains the learning process by making it more difficult to infer treatment quality from health outcomes.

4.3.1 Treatment Noise

While under treatment patients learn about their personal effectiveness of treatment. Each observation under treatment helps the patient determine and separate their personal effectiveness of treatment from the health signal noise. The greater the signal noise the harder it is for patients to discern between the true treatment effect and the signal noise. Signal noise ultimately slows the learning process of patients. Conditional on the patient's initial prior mean and set of observed health outcomes, an increase in the variance of the signal noise lowers the value of continuing treatment such that $\frac{dz_t|F_t(\cdot|\bar{h},\mu_0)}{d\sigma_\varepsilon^2} < 0$.[9]

Although signal noise lowers the value of treatment (conditional on the patients prior and set of prior outcomes) it does not necessarily lower adherence. This is because the impact of signal noise on adherence interacts with observed health outcomes. Consider the extreme case where the treatment effectiveness can be inferred from health outcomes immediately such that $\sigma_\varepsilon^2 = 0$. The fraction of responders for which treatment is valuable $(U(q,p) > U(q_s,p_s))$ adhere while the fraction of non-responders immediately leave. In the other extreme, when nothing is learned from health outcomes about treatment effectiveness, σ_ε is extremely large. In this case, little is learned from treatment experience and adherence depends strongly on a patient's prior beliefs (assumed to be the population average treatment effect). If the average treatment effect is positive then everyone adheres, otherwise no-one adheres. These extreme cases illustrate how the

[9] See Lemma 1 in Yao (2006) for details on the proof.

impact of signal noise on adherence interacts with the quality of treatment. In general, decreasing the noise variance produces an indeterminate effect on adherence behavior.

4.3.2 Heterogeneity in Treatment Quality

An inherent value of the treatment regime is the learning and option value of continuing treatment. The patient's outside alternative to adopt the standard of care allows patients to partake in the upside of treatment quality without the downside risk. The option value of treatment and hence the value of the treatment regime is increasing in the variance of treatment quality σ^2. Conditional on the patient's prior mean at time t, an increase in the variance of treatment quality increases adherence such that $\frac{dz_t|F_t(\cdot|\bar{h},\mu_0)}{d\sigma^2} > 0$.[10] However, an increase in variance of treatment quality also increases the dispersion of health outcomes which could potentially increase or decrease overall adherence. Just as with the signal noise, a change in the variance of treatment quality in general produces an ambiguous effect on treatment outcomes.

4.4: Education and Adherence

We consider when education affects the ability to learn, in particular the type of Bayesian learning about treatment value discussed here. The analog arguments apply to a patient who has specialized education in health or treatment related matters, so called "health-literacy", potentially by learning over time about a chronic condition. One interpretation of being more educated and informed is that beliefs are closer to the truth than when less educated. In our framework, more educated may be either interpreted as a closer prior mean (μ_0) to the true mean quality or a lower prior variance (σ_0^2) or noise of the outcome (σ_ε^2).

---

[10] See Theorem 1 in Yao (2006) for details on the proof.

Consider the scenario where education shifts the perceived effectiveness of the treatment. For example, educated people may believe less in holistic medicines but more in therapies tested in trials. Interpreted this way, education does not produce a monotonic effect on adherence but interacts with treatment effects. Education therefore will have a positive (negative) effect on adherence when the education increases (lowers) the perceived quality of the treatment regime.

An alternative interpretation of education is that it impacts the perceived distribution of treatment effects. Suppose more educated individuals have more precise knowledge of their own personal treatment effect (lower σ^2). Under this interpretation, education increases each patient's knowledge of the personal effectiveness. As discussed in earlier sections, a decrease in the patient's prior variance actually lowers the value of treatment at any time t conditional on the patient's perceived effectiveness (prior mean). In this sense education lowers the option value of treatment adherence which lowers adherence. An interesting empirical question regarding education effects is whether MDs adhere differently to treatments than those less educated. Our education effect implies that MDs may have stronger priors over the value of treatment. The option value of treatment is lower for MDs which implies MDs only initiate treatment programs they know will be effective.

Alternatively education could be interpreted as a shift in the distribution of signal noise. Education could lower variance of the idiosyncratic signal noise component of health outcomes. For example, educated individuals may be more likely to take the medication as prescribed by their physicians by following the correct dosages, timing and other factors (dietary restrictions, etc.). Conditional on the perceived effectiveness of treatment, a decrease in the variance of the signal noise increases the value of treatment. Under this interpretation, education helps patients

infer true treatment quality from signal noise.[11] In this sense, education produces potentially important welfare effects when examining ex-ante versus ex-post efficient adherence.

The discussed predicted interactions between education, observed treatment performance, and adherence can be tested with individual level outcomes data under the maintained assumption that priors correspond to the true population wide distribution of effects.

4.5: Real World vs Trial Effectiveness and Adherence

The so called "real world" effectiveness of a treatment after a product has been marketed may differ from the effectiveness of the product as found in a clinical trial prior to marketing. One major factor that may differ between a trial and a real world setting is adherence as it is often encouraged and monitored more extensively in trials. Consider a treatment with the trial effect corresponding to the average effectiveness μ above. Define the real world average effectiveness $\mu_R(t)$ at a given point in time as the one resulting from those still adhering to the treatment as in

$$\mu_R(t) = \int_{-\infty}^{\infty} q \frac{S(t|q)}{S(t)} dF(q)$$

In other words, real world effectiveness is the selected effectiveness that occurs conditional on the patients still on treatment.

Our analysis has implications for the difference between real world effectiveness and trial effectiveness. First, as the price p lowers adherence survival $S(t|q)$ for each quality level, price

[11] A patients prior variance of the personal effectiveness of treatment at time t is given by $\sigma_t^2 = \left(t\sigma_\varepsilon^{-2} + \sigma_0^{-2}\right)^{-1}$. A decrease in the variance of the signal noise strictly increases each patient's knowledge of their own personal treatment effectiveness.

differences between trials and real world treatment settings are important. In particular, if trials are free or subsidized to patients but treatments are priced when marketed, then adherence will clearly be lower in a real world setting. Clinical trial effectiveness will underestimate real world effectiveness. Consider when, regardless of health outcomes, adherence to the treatment regime in a randomized clinical trial occurs so that $S(t|q) = 1,\ \forall q, t$. In the real world setting, patients optimally non-adhere over time when the average experienced effectiveness of treatment falls below the cutoff value. Thus, as discussed, the real world conditional survivor function is increasing in personalized treatment effectiveness, $\frac{\partial S(t|q)}{\partial q} > 0$. Consequently, the real world average treatment effectiveness will be greater than the average trial treatment effectiveness. Treatments look worse in trials in compared to the real world when those performing worse do not leave trials to the same degree they do not adhere when the treatment is marketed.[12]

Third, real world effectiveness rises disproportionally with the trial effect. This is because when μ rises it affects both the distribution of qualities $F(q)$ as well as redistributes patients to adherence survival curves $S(t|q)$ that are higher. In other words, a better treatment does not only perform better but also make patients adhere more to the better performance.

Lastly, the difference between real world and trial effectiveness interacts with the true treatment effect. As patients stick to good treatments, the survival $S(t|q)$ goes to unity as the quality q rises. This therefore implies that real world effectiveness converges to trial effectiveness as the treatment improves

[12] There are other conditions in trials, such as blinding or more disperse beliefs about treatments, which may have opposing effects to their lower prices. The discussion here assumes that the price effect dominates, but an analog discussion applies if it does not.

$$\lim_{\mu \to \infty} \mu_R(t) - \mu = 0$$

In other words, as the treatment performs better so does adherence, resulting in real world effectiveness that gets closer to trial effectiveness.

4.6: Treatment Duration and Adherence

The hazard rate of non-adherence is defined as the fraction of remaining patients that quit the treatment in a given period. In our analysis, learning about the quality of the treatment takes place initially but eventually the patient learns its value with great precision. This implies that after a sufficient amount of treatment, the hazard rate into non-adherence goes to zero, that is, secondary non-adherence vanishes.

As treatment progresses a patient's observed average treatment effect converges to the true personalized quality of treatment, q. This is illustrated through the variance of the patient's posterior (eq. 1), as t increases the posterior variance converges to zero,$\lim_{t \to 0} \frac{1}{\frac{t}{\sigma_\varepsilon^2} + \frac{1}{\sigma^2}} = 0$. Once patients know the true treatment quality, they elect to adhere to the treatment if and only if $U(q, p) \geq U(q_s, p_s)$.

4.7: Comorbidities and Adherence

A patient undergoing a given treatment may be undertaking other treatments due to multiple diagnoses or comorbidities. There are three ways in which comorbidities may affect adherence in our analysis. First, the effectiveness of a treatment may depend critically on the patient's comorbidities and the associated treatments. For example, the effectiveness of one drug may be partially subdued or enhanced when taken in conjunction with another drug.

Secondly, comorbidities may make it harder for patients to infer treatment value from health outcomes. This is because when the patient is on several treatments due to comorbidities, the patient does not know whether it is the treatment itself, the comorbidities, or the treatment for comorbidities that may be causing a given health outcome. In other words, comorbidities raise the variance of the signal noise σ_{ε}^2. As discussed, a rise in the variance of the signal noise lowers the value of treatment but overall produces an indeterminate effect on adherence.

Lastly, comorbidities may affect adherence by making it marginally more taxing on a patient, both financially and mentally, to undertake multiple treatments for multiple morbidities. For example, it may be more taxing to remember when to take eight medicines rather than one. This would be reflected by a higher total price p in our analysis and clearly decreases adherence. The overall effect of comorbidities on adherence will be determined by whether the price effect dominates a potential positive effect of prolonging learning through a higher signal noise.

4.8: Providers and Adherence

Non-adherence is defined as not following recommended therapy from a third party provider such as say a doctor or nurse. Providers may affect adherence through determining the prior beliefs F_0 of the patient. This may occur through effective provider-patient communication about the benefits of the treatment and adherence to it. There are several ways in which providers may drive adherence through affecting the beliefs of the patient. The first is that the reimbursements of providers are tied to adherence and therefore providers undertake more or less effort in educating the patient about the value of adherence. For example, in a fee-for-service setting a provider may actually obtain larger reimbursements from poor adherence by the patient returning for subsequent care due to poor adherence. Or providers may monitor adherence more

closely and communicate with patients better when they are directly rewarded for adherence as has been the case for larger reimbursements when vaccination schedules are adhered to in some countries such as the UK.

A provider's impact on adherence depends on the quality and trustworthiness of the provider. If a well-known physician from a prestigious medical institution prescribes a treatment, the patient may believe in the treatment more than if a resident or nurse from a community hospital prescribed it. Better doctors may have better patients because their patients place more weight on the opinion/ beliefs conveyed by their doctor.

If providers educate patients about the value of treatment, there are similarities between provider and education effects on adherence. One way in which providers may educate their patients is by informing them of the population wide average treatment effects which will impact a patient's prior mean and variance. Under such an interpretation, provider effects on adherence will mimic education effects on the adherence.

Similar to the case of education, providers will not have an effect on the long run belief of the patient. Ultimately the patient will come to learn whether the treatment works for him or not which will drive adherence. This is reflected in that posterior beliefs are determined more and more by the experienced health outcomes over time regardless of the patient's prior belief in the treatment. However, by affecting prior beliefs the provider may prevent patients from exiting treatment due to initially poor performance, something we discuss later in the normative analysis.

Section 5: Normative Implications for Efficient Adherence

In this section we discuss the efficiency implications of non-adherence.[13] Inefficient adherence transpires as the direct result of heterogeneous and unknown personalized treatment effects. The process in which individuals learn about their own treatment creates the potential for both under and over-adherence.

5.1 Ex-ante vs Ex-post Efficient Adherence

Ex-ante efficient behavior occurs if an individual cannot be made better off given their individual information at a given point in time. By definition, stopping behavior being individually ex-ante optimal implies that adherence is ex-ante efficient unless there are external effects (we discuss such issues in the conclusion). *Ex-post efficient* behavior occurs when only those who actuality value the treatment adhere to it. Let q^* be threshold level of health or treatment quality (the same thing ex-post) which makes the patient is indifferent between the treatment and the standard of care

$$U(q^*, p) = U(h_s, p_s)$$

Naturally, the reservation level of health q^* increasing in the price of the treatment and the health of the standard of care but decreasing in the price of the standard of care price; $\frac{\partial q^*}{\partial p} > 0$, $\frac{\partial q^*}{\partial q_s} > 0$, and $\frac{\partial q^*}{\partial p_s} < 0$.

[13] Both medical and economic discussions of adherence often state that patients do not adhere enough, although there is no explicit criteria discussed defining whom and why a patient should adhere. In some sense, our theory suggests an explanation of this normative claim by third party bystanders about under-consumption of patients; the selection effect inherent in learning means that those that adhere do better than those who do not adhere. With the inherent upward bias in adherence effects under optimal learning, it maybe ill-advised to argue everyone should adhere.

It is ex-post efficient for patients with personalized treatment effects above this reservation level $q \geq q^*$ to be on treatment. Therefore, there are two types of ex-post inefficiencies. The first inefficiency is under-adherence by responders for whom treatment is valuable. Even though treatment is valuable for a fraction $1 - F(q^*)$ of patients, some of those patients will stop treatment because of incorrect inferences about treatment value. The second type of inefficiency is over-adherence by non-responders for whom treatment is not valuable. Some of the non-responsive patients, of size $F(q^*)$, will initially adhere to the treatment before learning that it is not valuable.

Figure 1 illustrates the general pattern of ex-post inefficient adherence. The sold gray line reflects the proportion of individuals at each period that adhere for whom the treatment is valuable and is reflected by the survival curve $S^E(t)$ indicating efficient adherence. It is the survival function for which the new treatment is optimal, i.e. $q > q^*$. Under-adherence by this group is reflected by the fact that the solid gray line is not equal to one; $S^E(t) < 1, \forall t > 0$. This occurs because some of the patients experiencing poor initial performance on the treatment leave even though it is in fact valuable.

The dotted line reflects the proportion of individuals at each period that adhere for whom the treatment is not valuable (i.e. non-responders) as reflected by the survival curve $S^I(\cdot)$ representing inefficient adherence. It is the survival function conditional on $q < q^*$. By our previous discussion that argued that true treatment quality raises adherence, it will be the case that $S^I(t) < S^E(t)$, $\forall t$. Over-adherence by the non-responders is reflected by the fact that their survival curve is positive. This occurs because some non-responders will adhere due to the option value of treatment and/ or because by luck they experienced good initial performance on

the treatment. However, sooner or later all of them will learn that the treatment is invaluable so that no one adheres; $\lim_{t\to\infty} S^I(t) = 0$.

FIGURE 1: ADHERENCE BEHAVIOR BY RESPONDERS AND NON-RESPONDERS

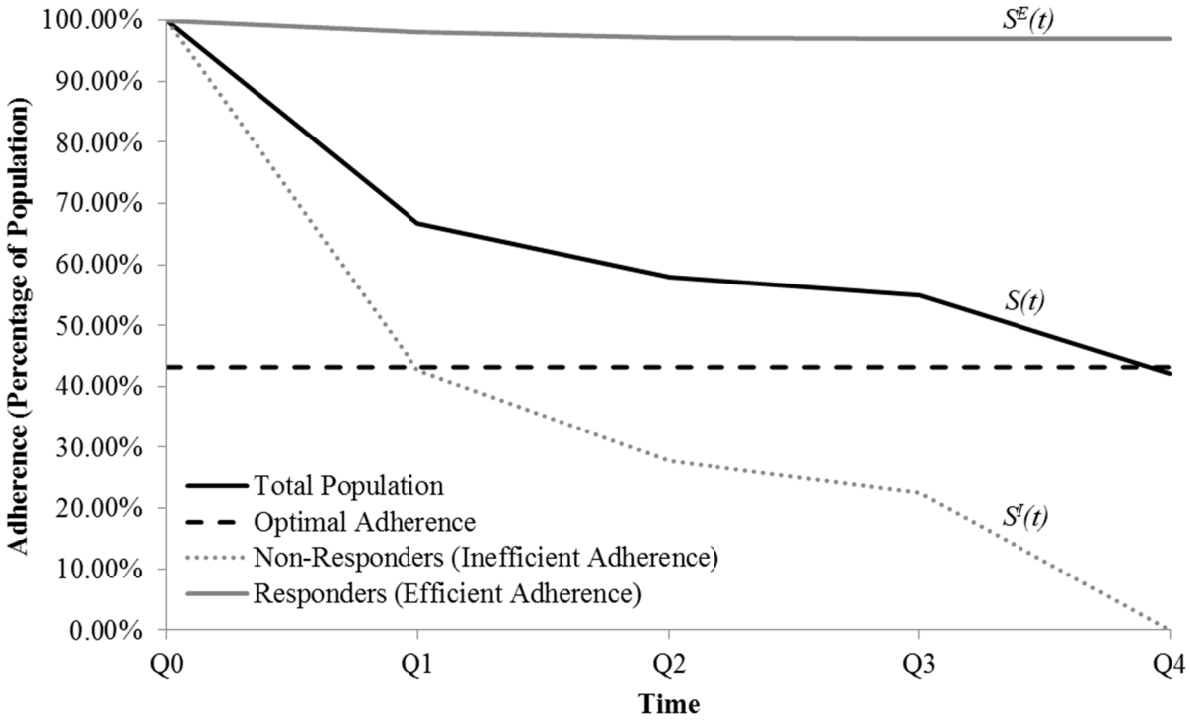

Notes: Figure 1 illustrates the adherence survival curves. The non-responders survival function, $S^I(t)$, illustrates adherence for those individuals for whom treatment is not valuable/inefficient. Similarly the responders survival function, $S^E(t)$, illustrates adherence for those individuals for whom treatment is valuable/efficient. Optimal adherence reflects the percentage of the population for which treatment is valuable/ efficient.
The survival curves are calculated according to the parameter values in Table 1 in the proceeding calibration section. For computational ease we assume that the true quality of treatment is revealed to adhering patients after one year of treatment.

The overall solid survival curve is the mixture of the two conditional survival functions with the mixture weights given by the fraction of responders and non-responders:

$$S(t) = F(q^*)S^I(t) + (1 - F(q^*))S^E(t)$$

The sold dotted line is the fraction of patients $1 - F(q^*)$ for whom the treatment is valuable, which is 43.16% in Figure 1. It follows that for the overall population there will initially be both under-adherence for responders and over-adherence for non-responders in the short run. However, in the long run there will always be under-adherence because there are patients who drop out and will never find it optimal to re-adhere; S^E of true responders inefficiently remains below unity. The survival S^I of true non-responders efficiently goes to zero as non-responders learn that the treatment is not worthwhile for them.

This previous discussion concerned the inefficiency in quantities, that is, who is on the treatment or not compared to who should be. The monetary value lost from under and over adherence results from how much the foregone optimal therapy is valued. Let the reservation price for an individual of type q be denoted $r(q)$ and is defined as

$$U(q, p - r(q)) = U(q_s, p_s)$$

Thus the sign of $r(q)$ reflects whether the treatment is truly valued relative to the standard of care and can be interpreted as the discount required to induce an individual of type q into treatment. It then follows that $r'(q) < 0$ and that for non-responders $r(q) \geq 0$ (i.e. $q \leq q^*$). The flow of welfare loss at time t can then be written as

$$\text{Welfare Loss at t} = W_O(t) + W_U(t) = \int_{-\infty}^{q^*} r(q)S(t|q)dF(q) + \int_{q^*}^{\infty} -r(q)\left[1 - S(t|q)\right]dF(q)$$

The first term $W_O(t)$ is the loss in welfare at time t from over-adherence; those who do not value the treatment but still adhere to it. The second term $W_U(t)$ is the loss in welfare at time t from under-adherence; those who value the treatment but stopped adhering. Figure 2 illustrates the overall cumulative effect of these two types of inefficiencies when the flows are aggregated up and weighted over time.

FIGURE 2: EFFICIENT AND INEFFICIENT ADHERENCE SURVIVAL CURVES

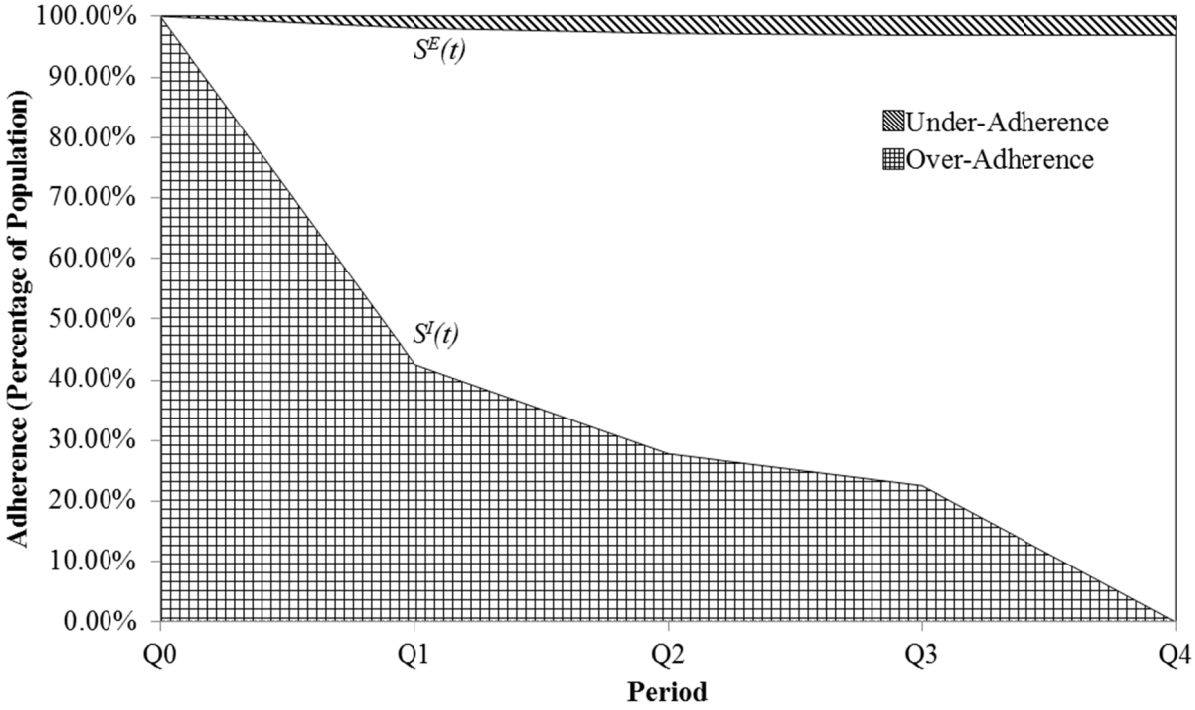

Notes: Figure 2 illustrates the simulated survival curves. The non-responders survival function, $S^I(t)$, illustrates adherence for those individuals for whom treatment is not valuable/inefficient. Similarly the responders survival function, $S^E(t)$, illustrates adherence for those individuals for whom treatment is valuable/efficient.
The survival curves are calculated according to the parameter values in Table 1 in the proceeding calibration section

The area above the responder's survival function (S^E) represents the extent of under-adherence: the fraction of patients that are not adhering but should not be. The area below the non-responders survival function (S^I) represents over-adherence: fraction of patients that are adhering but should not be. The present value of the total welfare loss is the discounted value of the loss from both forms of inefficiencies

$$W = \int_0^\infty \beta^t \left[W_O(t) + W_U(t)\right] dt$$

The important aspect of this overall welfare loss is that over-adherence is front-loaded while under-adherence that is back-loaded. Therefore, in present value terms over-adherence often matters more than under-adherence. This makes common non-explicit arguments about the loss from under-adherence clash with the present value of welfare effects.

5.2 Welfare Effects of Common Adherence Interventions

The ex-post inefficiencies in quantities or dollars discussed above are affected in non-obvious ways by interventions that aim to raise adherence. Consider when the adherence intervention represented by the parameter θ raises survivals of both responders and non-responders. For example, consider an education campaign that changes treatment beliefs and increases overall adherence. This implies that over-adherence is increased and under-adherence decreased; $\frac{dW_0}{d\theta} > 0$ and $\frac{dW_U}{d\theta} < 0$. The effect on the overall ex-post welfare is given by

$$\frac{dW}{d\theta} = \int_0^\infty \beta^t \left[\frac{dW_O(t)}{d\theta} + \frac{dW_U(t)}{d\theta} \right] dt$$

Therefore, any intervention that affects adherence behavior by the two groups *symmetrically* will have indeterminate effects on ex-post efficiency. To illustrate, consider the indeterminate welfare effects of lowering co-pays to raise adherence. A decrease in the price p raises adherence of both groups, S^E and S^I. Therefore, a cut in co-pays raises efficiency for responders, by making them adhere more when they should, but lowers efficiency for non-responders, by making them also adhere more when they should not.[14]

[14]Learning about treatment effects alters the standard tradeoff between risk and incentives inherent in discussions of optimal copay levels. This is true for both the case of a single service or good (Pauly 1968 and Zeckhauser 1970) as well as insurance for multiple goods and services[14] (Goldman and Philipson 2007).

In the long run, we know that the intervention effect on over-adherence must go to zero since over-adherence tends to zero as individuals learn about their personalized treatment effect. Therefore, in the limit the welfare effects of an intervention raising adherence must lower inefficient adherence and the associated flow costs

$$\lim_{t\to 0} \frac{dW_O(t)}{d\theta} = 0 \implies \lim_{t\to 0} \frac{dW(t)}{d\theta} < 0$$

Whether or not the intervention lowers the welfare loss depends on if the future benefit of decreased under-adherence more than offsets the short term costs of increased over-adherence. Determining the effect of adherence interventions on overall welfare is thus difficult from data on adherence and health outcomes alone and would likely require additional identifying restrictions. We provide an illustration of inferring these welfare effects from a structural model in the proceeding section.

Education and memory-enhancing interventions are often used to attempt to raise adherence. Ex-post efficiency may be improved if the patient is poorly educated and has the wrong treatment beliefs, even though ex-ante efficiency effects are more difficult to assess when changes in beliefs change preferences. In an ex-post sense, patient education may improve efficiency if patients are not fully aware of the benefits of treatment or if patients are too optimistic about treatment value. In those cases, there may be an efficiency role for education, bringing priors closer to the true treatment effect, in raising ex-post efficient adherence.

However, the patient's education or the type of provider he sees will not affect long run over-adherence by non-responders. This is because their adherence survival will go to zero in the long run because health outcomes dominate treatment beliefs in the long run. In other words, the weighting $(\omega, 1-\omega)$ of experience and priors in the posterior puts more and more weight on

experienced outcomes relative to prior beliefs. Neither the amount of education nor the type of doctor can make the patient stay on a treatment that the patient knowingly does not value.

Section 6: Calibrating Adherence Inefficiencies: An Illustration for Simvastatin (Zocor)

In this section we calibrate our model of non-adherence to assess the welfare losses induced by inefficient adherence. We consider adherence associated with cholesterol lowering treatments taken by adult males. More specifically, we calibrate our model of adherence to 58 year old males' adherence to the drug simvastatin (Zocor) as a cholesterol lowering treatment regime. Our main calibration result is that even though a majority of these patients do not adhere, the welfare loss of over-adherence dominates that of under-adherence. In particular, the present value of the per-capita loss due to over-adherence to simvastatin is over 80% larger than the loss due to under-adherence ($58 vs $7).

We interpret simvastatin (Zocor) as the unknown treatment while the alternative known standard of care is not taking any treatments. Our interpretation assumes that the sole objective of the treatment is lowering low-density lipoprotein cholesterol (LDL-C). The per period (quarterly) benefit of simvastatin is h_t represents the patient's percentage point decline in LDL-C levels relative to their initial baseline levels. The percentage point decline in LDL-C levels of a patient in a given period reflects the true personalized treatment effect plus some idiosyncratic shock according to

$$h_t = q + \varepsilon_t$$

where $q \sim N(\mu,\sigma)$ and $\varepsilon \sim N(0, \sigma_\varepsilon)$. Patients observe their cholesterol levels and update their adherence decision on a quarterly basis. For computational ease we further assume that all

patients initiate treatment and learn their true value of treatment after one year of treatment. Therefore, patients continue with the treatment after a year (hazard rate goes to zero) if and only if they are true responders i.e. ($U(q,p) > U(q_s,p_s) = 0$) for the remainder of their lives.

Using readily available clinical trial data, we calibrate the model for 58 year old males. We therefore assume that the each patient expects to live the average life expectancy of 23 years without treatment but longer if responding to the simvastatin treatment.[15]

This calibration requires knowledge of the distribution of treatment effects, (μ, σ^2), the distribution of signal noise (σ_ε^2), the costs of treatment (p), and the utility parameters (β, γ). Clinical trial data in principle can provide information on the treatment quality and noise parameters. The treatment mean and variance (μ, σ^2), is often directly reported from such trials and individual longitudinal data can be used to estimate the noise distribution (σ_ε^2). For the calibration exercise here, we use clinical trial data on the distribution of effectiveness of simvastatin from (Bays et al. 2004). Table 1 summarizes the parameters values used in calibrating the model. On average, the simvastatin treatment therapy in the Bays et al. study lowered LDL-C levels by 37.00% over a quarter relative to a placebo.[16] In the context of the model discussed in the previous section, this implies $E[h] = E[q] = \mu_0 = 37.00\%$ and $h_s = q_s = 0.00\%$. We use simvastatin treatment cost estimates from Hoadley et al. (2012). In particular, using Medicare Part D data, Hoadely et al (2012). find that the median out of pocket cost paid by users for a one quarter supply of branded simvastatin (Zocor) was $231.25.[17,18]

---

[15] We calculate life expectancy according to the 2009 CDC National Vital Statistics Report and the Social Security Administration.

[16] Patients studied in Bays et al (2004) received dosages of 10-80 mg respectively of simvastatin per day.

[17] Hoadely et al. (2012) find that the median 30-day out of pocket cost for branded Zocor was $71 in 2008. We convert their cost estimates into the cost of a one quarter supply in 2014 by scaling the cost by 3.257 to account for

TABLE 1: PARAMETERS VALUES

Parameter	Value
Mean Effectiveness of Simvastatin (Zocor) Therapy (μ_0)‡	37.00% per quarter
SD of Effectiveness of Simvastatin (Zocor) Therapy (σ_0)‡	14.80% per quarter
Cost of Simvastatin (Zocor) Therapy (p)	$231.20 per quarter
Signal to Noise Ratio ($\sigma_0^2/\sigma_\varepsilon^2$)	2.00
Discount Factor (β)	0.90
Health Consumption Trade-Off (γ)	0.17% per dollar

‡ These parameter values are from the clinical study Bays et al. 2004. The study looks at the effect of a simvastatin (Zocor). Effectiveness measures the percentage point drop in low density lipoprotein cholesterol (LDL-C) over one quarter relative to the initial baseline level.
We calculate the cost of Zocor using the observed median out of pocket cost as calculated in Hoadley et. al (2012). See footnote 17 to see how the cost estimate is adjusted for inflation and dosage.
The discount factor and signal to noise ratio are calibrated to fit the empirical survival function for statin adherence estimated in Yeaw et al. (2009).
The health consumption trade-off parameter represents a patient's willingness to pay to lower their cholesterol in percentage points. We calculate the health consumption trade-off parameter as described in the text using existing value of a statistical life year (VSLY) estimates and the longevity benefits of simvastatin.

The health consumption trade-off parameter γ represents a patient's willingness to pay to lower his cholesterol for one period (one quarter). We calibrate the health consumption trade-off parameter γ using data on the longevity gains from simvastatin and existing value of statistical life year (VSLY) estimates. Based on the results from the Scandinavian Simvastatin Survival Study (S4)[19], Jönsson et. al (1996) find that simvastatin treatment saved an estimated 0.377 undiscounted life years. These estimated longevity effects are in line with the results from the

---

the quantity and inflation. We account for inflation according to the BLS inflation calculator [http://www.bls.gov/data/inflation_calculator.htm]. We find similar cost estimates using CVS Pharmaceutical data from GoodRx.com . GoodRx reports estimated cash price of a one month dosage (taken daily) of 20mg simvastatin Zocor at CVS Pharmacy is $38. Assuming that each patient receives 30mg of simvastatin daily implies the cost of a one quarter dosage is then $171.

[19] Patients were given 20-40mg of simvastatin daily over a roughly five year period (5.4 years on average). Over the whole course of the study, simvastatin lowered LDL-C levels by 35% on average (Pederson et al. 1994). These findings are similar to those in Bays et. al 2004 study.

Heart Protection Study Collaborative Group (2006) study. This implies that taking simvastatin for an average 58 year old male increases his life expectancy from roughly 81 to 81.377 years.[20,21] Existing VSLY estimates provide a valuation of the additional 0.377 gain in life years at the age of 81. We use the VSLY estimates from Murphy and Topel (2006) who estimate the value of a life year for an 81 year old at roughly $230,000 in 2014.[22] Under the assumption that the only benefit of simvastatin is increased longevity, we equate the discounted stream of health benefits $\frac{1}{\gamma}\mu$ (expressed in dollars) with the longevity benefits.

$$\sum_{t=0}^{23\times 4} \beta^{0.25t} \frac{1}{\gamma}\mu = \beta^{23} VSLY_{81} \times 0.377 \tag{2}$$

Given our parameter estimates of μ, $VSLY_{81}$, and β, we solve for the health consumption trade-off parameter in equation (1) finding that $\gamma = 0.17\frac{\%}{\$}$. In other words, patients are willing to pay one dollar to lower their LDL-C levels by a bit more than a sixth of a percentage point, 0.17%, per quarter.

The remaining parameters in the model to be calibrated are the discount factor β and the signal to noise ratio $\sigma_0/\sigma_\varepsilon$. We calibrate β and $\sigma_0/\sigma_\varepsilon$ to match observed adherence patterns for statin usage post-approval. Using pharmaceutical claims data, Yeaw et al. (2009) estimate the adherence survival function we discussed for statin adherence.[23] We calibrate β and $\sigma_0/\sigma_\varepsilon$ to minimize squared differences between the calibrated and empirical adherence survival function

[20] The average age in S4 study for males was 58 .1 years old (Pederson et al. 1994). The average age in the Bays et al. (2004) study was 56 years old.
[21] We calculate life expectancy according to the 2009 CDC National Vital Statistics Report and the Social Security Administration.
[22] See Figure 2(b) in Murphy and Topel (2006). Since Murphy and Topel's value of life year estimates are expressed in USD 2000 we adjust them by a factor of 1.38 to express the estimate in USD 2014 according to the BLS [http://www.bls.gov/data/inflation_calculator.htm].
[23] Although Yeaw et al. examine adherence to all statins, not just simvastatin, studies have shown that any of the statins available in the US are effective for moderate (up to 35%) LDL-C cholesterol reductions (Smith et al. 2009).

at each quarter for the first year.[24] The calibrated health discount factor is 0.90 which is line with the estimates from Moore and Viscusi (1988) and Viscusi and Moore (1989).[25] The calibrated signal to noise ratio is 2.00. This implies that the standard deviation of patient specific quarter-to-quarter to cholesterol levels of patients on simvastatin is $\sigma_\varepsilon = 7.40\%$. Note that in principle, the signal noise to ratio could be estimated using longitudinal clinical trial data on health outcomes. When such data is available, one would not need to observe adherence data in order to calibrate the needed parameters, thereby allowing for out of sample predictions about future adherence behavior from trial data obtained before marketing of the product.

Figure 3 below (and Figures 1 and 2 in the preceding section) display the survival function from the calibrated model. We calculate the survival function by simulating the model with 10 million hypothetical patients with the parameter values displayed in Table 1. The solid black line reflects the survival function corresponding to the calibrated model while the solid gray line reflects the observed adherence survival function as reported in Yeaw et al. (2009). The calibrated survival function closely mirrors the observed empirical survival function, exhibiting a correlation of 0.94.

[24] Empirically, we find the values of β and $\sigma_0/\sigma_\varepsilon$ by implementing a grid search over the parameter space $\beta \in \{0.9, 0.95, 0.99\}$ and $\sigma_\varepsilon/\sigma_0 \in \{0.50, 0.75, 1.00, 1.25, 1.50\}$.

[25] See Moore and Viscusi (1990) for further discussion on estimating discount rates for health outcomes.

FIGURE 3: CALIBRATED ADHERENCE SURVIVAL FUNCTION

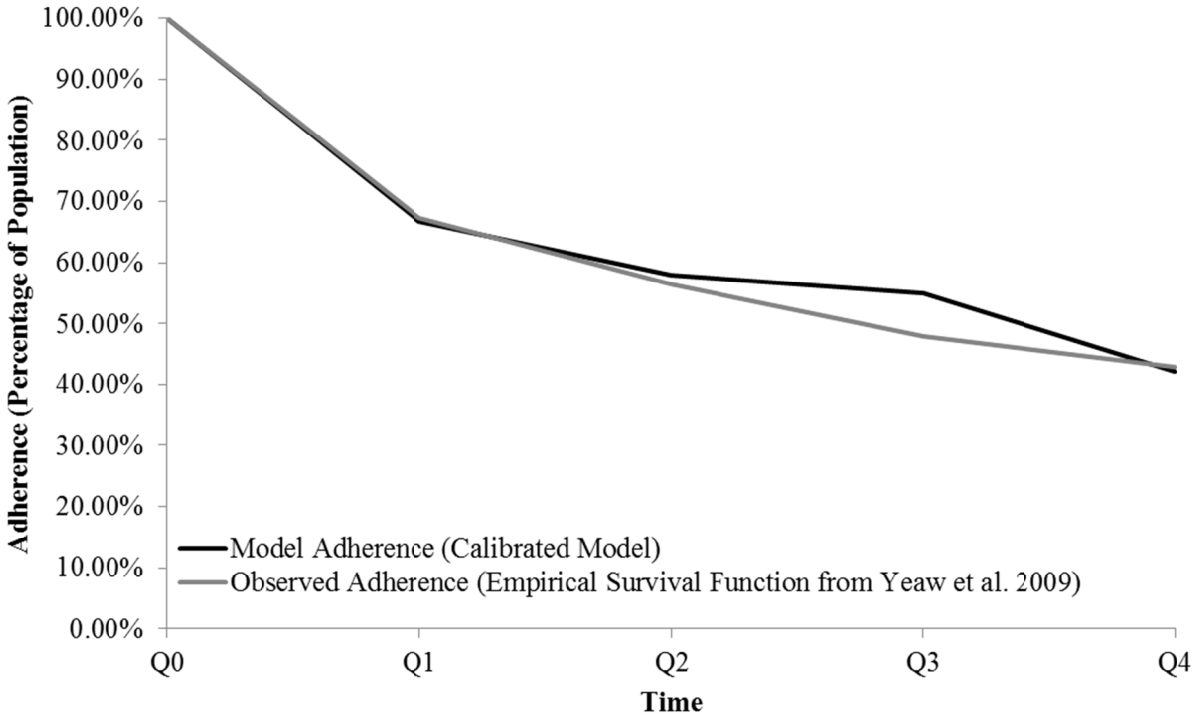

Notes: Figure 3 illustrates the calibrated survival curve corresponding to the parameter values in Table 1 and the empirical survival curve estimated in Yeaw et al. (2009).

Both the observed and calibrated adherence to simvastatin is relatively low displaying that the majority of patients do not adhere to the treatment in the long run. In particular, the parameter values imply that the simvastatin treatment is adhered to by only 42.91% of those prescribed. The large degree of non-adherence is driven by the discounting of the future modest longevity effect of the treatment. The calibrated discount factor of 0.90, which as discussed is in line with previous estimates (Moore and Viscusi 1990), implies that patients are not willing to trade off current consumption for health benefits occurring in the distant future. In addition, the more the patients discount the future, the less weight patients place on the discussed option value of treatment.

The model calibration allows us quantify inefficient adherence and disentangle the inefficiencies driven by over-adherence vs under-adherence. Figure 2 in the previous section illustrates the calibrated survival functions for non-responding (S^I) and responding (S^E) simvastatin users using the parameter values of Table 1. These two survival functions suggest that over-adherence may be more problematic than under-adherence for simvastatin. About 42% of true non-responders still take simvastatin after a one period. However, as discussed, over-adherence vanishes as patients learn about the quality of the drug while the under-adherence problem persists. In the long run 3.1% of true responders under-adhere to the treatment.

Using the calibrated model, we can calculate the welfare loss associated with inefficient adherence. More precisely, we are able to monetize the welfare costs of inefficient adherence using the health consumption trade-off utility parameter γ. Consider a patient adhering to the simvastatin treatment. The simvastatin treatment is inefficient if and only if $q - \gamma p < 0$ or equivalently $r(q) \geq 0$ The associated cost of over adherence is equal to the cost treatment minus dollarized cholesterol effect.

$$\text{Quarterly Cost of Over Adherence} = p - \frac{1}{\gamma} q$$

By definition, the cost of over adherence is strictly positive for those individuals who are inefficiently over-adhering. An analogous expression applies to those who truly value the new treatment but stop adhering.

$$\text{Quarterly Cost of Under Adherence} = \frac{1}{\gamma} q - p$$

Figure 4 below displays the model implied per patient welfare loss from over and under adherence.

FIGURE 4: WELFARE LOSS OF EX-POST INEFFICIENT ADHERENCE TO SIMVASTATIN

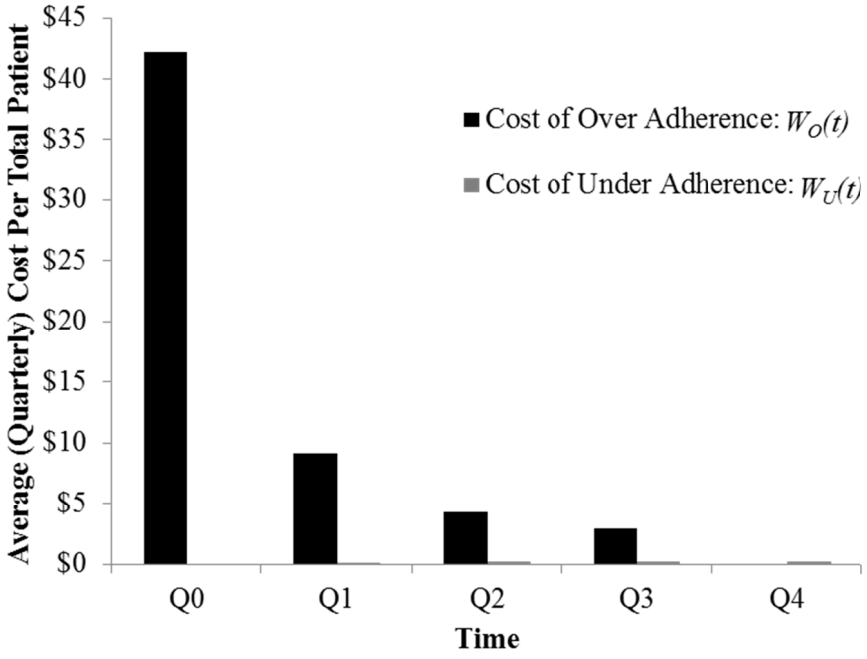

Notes: Figure 4 illustrates the costs of over and under adherence corresponding to the parameters in Table 1. The costs of over-adherence at each period is calculated as sum of $p - \frac{1}{\gamma}q_i$ across all adhering non-responders normalized by the total number of patients. The cost of under adherence is calculated as the sum of $\frac{1}{\gamma}q_i - p$ across all non-adhering responders normalized by the total number of patients.

The black bars in Figure 4 represent the average welfare loss per patient in a given quarter stemming from over-adherence by patients for whom the simvastatin treatment is not valuable. The gray bars represent the average loss per-patient from under-adherence by patients for whom the simvastatin treatment is indeed valuable. The cost of over adherence goes to zero as patients learn that the treatment is not valuable. Since individuals never re-enter treatment, the costs of under-adherence rises over time but converges to a steady state level in perpetuity as those who value the treatment stay on. The initial cost of over-adherence is $42.18 per patient quarter but declines to zero as non-responders drop out of treatment. Conversely, the cost of

under-adherence is initially zero as everyone exhibits primary adherence but is $0.21 per patient quarter in the long run. In present value terms, the total cost of under-adherence is $6.52 per patient and for over-adherence is $57.87 per patient.[26] The present value costs of over and under adherence are closer than quarterly levels because of the long-run nature of under-adherence. The total per capita cost of inefficient adherence (combining over and under adherence) is $64.39. To put these numbers in perspective, 94.1m simvastatin prescriptions were dispensed in the US in 2010 and total expenditure on lipid regulators was $18.7bn (IMS Health 2011). The potential aggregate costs of inefficient adherence are thus on the order of billions of dollars.

Section 7: Multiple Treatments and Partial Adherence on a Single Treatment

The previous analysis assumed the price and quality of the standard of care was known. This section generalizes the discussion to the case when the quality of more than one treatment is unknown to the patient. More precisely, we extend our previous adherence decision problem with two treatment alternatives, one uncertain treatment and one certain standard of care, to the more general setting with K uncertain treatment alternatives. One can think of the K treatments as completely different treatments or alternatively as different levels of adherence with the same treatment. For many treatments, patients may not fully discontinue a treatment but partially adhere to treatment each period, which thus may result in inferences about how partial adherence levels may affect health. Such partial adherence may be the case for many chronic conditions where patients do not undertake 100% of prescribed care but may skip or miss doses, e.g. insulin injections by diabetics.

[26] This is calculated using an annual discount factor of 0.90 and assuming that non-adhering patients live 23 years.

7.1: Adherence with Multiple Uncertain Treatments

Multiple uncertain treatments involve a so called multi-armed bandit problem in statistical decision theory. More precisely, following the previous set up, we assume that each of the K treatment alternatives produces a personalized health benefit which is a function of the treatment quality, q^k, and an idiosyncratic noise term ε_t^k,

$$h_t^k = q^k + \varepsilon_t^k,\ k = 1, 2...K$$

The personalized quality of treatment k is distributed i.i.d. across individuals from the distribution $q^k \sim F^k(\cdot)$. As before, each treatment generates utility $U(h_t^k, p_k) = h_t^k - \gamma p_k$.

Patients' prior belief over the quality at time t for treatment k id denoted $F_t^k(\cdot|\vec{h}_t^k)$ where $\vec{h}_t^k$ is the history of experienced personal health outcomes on treatment k. This formulation assumes that treatment qualities for a patient are distributed conditionally independently across the treatment alternatives.[27] Further we assume that each patients initial prior reflects the true distribution of treatment heterogeneity, $F_0^k = F^k$. Under the maintained assumption that the prior and shocks are normally distributed, $F_0^k(q) \sim\ N(\mu_0^k, \sigma_{k0}^2)$ and $\varepsilon_t \sim N(0, \sigma_{k\varepsilon}^2)$, standard normality results imply that a patient's posterior distribution is optimally updated according to the previous equation (1) the single treatment case.

The patient's adherence problem now involves selecting the optimal treatment regime among the K alternatives each period. The value function of a patient adhering to treatment option k at time t is

[27] See Pandey et al. (2007), Rusmevichientong and Tsitsiklis (2010) and Dickstein (2014) for a discussion of multi-armed bandit problems with correlated arms.

$$V(h_t^k, \mathbf{H}_t, \vec{F}_0) = U(h_t^k, p_k) + \beta \max \left\{ E[V(h_{t+1}^1, \mathbf{H}_{t+1}, \vec{F}_0)|\vec{F}_t], ..., E[V(h_{t+1}^k, \mathbf{H}_{t+1}, \vec{F}_0)|\vec{F}_t], \right\}$$

Here, $\mathbf{H}_t$,represents the matrix of outcomes across the K different treatments and the vectors $\vec{F}_0$ and $\vec{F}_t$ represent the patient's prior and posterior distributions over the K different treatment alternatives. The previous discussion with an uncertain treatment ($k = 1$) and a certain standard of care ($k = 2$) corresponds to when $\sigma_2^2 = 0$. As before, patients use their posterior treatments beliefs to optimally adhere to each treatment.

Patients optimally adhered in the case of two treatments by following a stopping rule that determined adherence behavior on the uncertain treatment. With multiple unknown treatments, the optimal adherence rule generalizes to selecting the treatment alternative with the highest Gittins index (Gittins and Jones 1974, Gittins and Jones 1979, Gittins et al. 2011). The Gittins index, I_t^k, for a particular treatment k at time t corresponds to the level of utility generated by some hypothetical known standard of care for which the patient is indifferent between treatment and the standard of care in the simple two treatment alternative case.[28]

$$\frac{I_t^k}{1-\beta} = E\left[U(h_{t+1}^k, p_k) + \beta \max \left\{ E[V(h_{t+2}^k, \vec{h}_{t+2}^k, F_0)|F_{t+1}], \frac{I_t^k}{1-\beta} \right\} |F_t \right]$$

The Gittins Index Theorem (Gittins and Jones 1974, Gittins and Jones 1979) shows that in each period, patients optimally adhere by selecting the treatment alternative with the highest Gittins index at that time. The Gittin's Index Theorem essentially reduces the K multi-armed bandit problem into a set of K single-armed bandit problems.

[28] See Powell and Ryzhov (2012) for a full discussion of Gittins indices.

7.2: Implications for Non-Adherence among Multiple Treatments

Computing Gittins indices correspond directly to the single armed bandit problem and optimal stopping rules described in Section 3 which govern optimal adherence. Furthermore, since Gittins indices in the multiple treatment framework are computed using the simple two treatment framework (with one known treatment alternative), virtually all comparative statics discussed in Section 4 generalize to the multiple treatment setting. For example, the Gittins index for a particular treatment k is increasing in the perceived quality of treatment (μ_t^k) while decreasing in the cost of treatment (p_k). Similarly, conditional on the perceived quality of treatment, the Gittins index for an alternative k is increasing in the variance of treatment quality (σ_{k0}^2) while decreasing in the variance of the signal noise ($\sigma_{k\varepsilon}^2$).

7.3: Implications Partial Adherence on a Single Treatment

The multiple-treatment framework allows one to assess behavior involving partial adherence. Different levels of adherence, such as fractions of prescribed medications taken, can be thought of as separate treatments in the multiple treatment framework. Consider a patient facing the option of fully adhering versus partially adhering to a treatment regime. On one hand, fully adhering to the treatment regime likely generates superior and less noisy health outcomes relative to partial adherence. Both of these attributes (higher mean and lower signal noise variance) make full adherence an attractive alternative relative to partial adherence. On the other hand, partial adherence is likely at substantially lower cost than full adherence, whether in direct treatment costs or time costs of compliance. Because of its lower cost and greater variance of treatment effectiveness, patients may find it optimal to partially rather than fully adhere to

treatment. The general point is that whatever effects that one believes are true for different levels of adherence can be viewed as multiple treatments with different health outcomes and costs.

Section 8: Concluding Remarks and Future Research

Little explicit positive and normative analysis exists in health economics on the dynamic demand behavior implicit in non-adherence which is often associated with uninformed patients. We analyzed the implications for adherence behavior stemming from patients learning about personalized treatment value. We stressed that although providers may be more informed about the population-wide effects of treatments, patients may be more informed of their own value of care in terms of how they trade off effectiveness, side-effects, costs of care and compliance. We derived the optimal stopping problem which corresponds to non-adherence under personalized patient learning and characterized its observable determinants. We also derived the normative implications resulting from such non-adherence and calibrated the welfare losses implied for the cholesterol reducing therapy simvastatin (Zocor). The calibration results suggest that losses due to over-adherence are over 80% larger than losses from under-adherence even though only 43% of patients adhered to the therapy.

We conclude by discussing some of the shortcomings of the analysis and issues that may be usefully considered in future theoretical or empirical research on adherence.

External Effects and Non-Adherence:

We only considered adherence from the private choice perspective of the patient. However, privately optimal adherence may not be socially optimal when adherence behavior confers external effects. For example, adherence to treatments for infectious diseases such as TB

may involve positive externalities and thus may be inefficiently low when non-infected individuals benefit from adherence by infected patients. For classes of drugs like antibiotics or antiretrovirals, there is an additional issue of the negative externality that non-adherence imposes on everyone else due to population resistance to the treatment. Or external effects may operate through insurance premiums when non-adherence raises the total cost of care through cost-offsets (Goldman and Philipson 2007, Chandra et al 2010). Pigouvian subsidies to stimulate adherence under positive external effects may then be relevant and may be implemented through lower copays or other methods that raise adherence. More careful analysis of the role of adherence programs is needed in the context of external effects.

Selection and the Effects of Adherence on Health

Medical studies stress the importance of adherence because of the positive impacts on a patient's health. For example, many analysts think patients need to be better educated about treatments for breast cancer given that compliance is poor but the health benefits seem substantial. However, our analysis directly implies that those adhering perform better than those that do not, and thus the adherence effects are over-estimated when optimal stopping occurs due to poor performance. The basic view of the medical community that patients under-consume care needs to be evaluated not from the average experience but from the patient specific experience. This selection also affects the optimal targeting of adherence interventions; low levels of adherence may reveal preferences that imply small effects for adherence interventions.

Insurance Design and Adherence

Future explicit analysis may usefully consider optimal insurance design when patient learning drives adherence. One direct way in which insurance may affect adherence is through co-pays as predicted. But other indirect ways include affecting the beliefs of patients through communication, or initial or temporary copay rebates to prevent under-adherence by those who benefit from care. There is an interesting health policy literature on so called value based insurance design (VBID) which discusses copay design. However, to our knowledge the optimality criteria that determine whether copays should be set high or low (i.e. the *definition* of V in VBID) are not explicitly discussed in that policy literature. This is in contrast to economic analysis of the value of insurance designs where the value V is defined as economic efficiency (see Pauly 1968 and Zeckhauser 1970 for co-pay design for single services and Goldman and Philipson 2007 for multiple services). Without an explicit definition of the value V in VBID, it is impossible to determine whether high or low copays are good or bad in a normative sense and may thus lead to inefficient recommendations. For example, in our framework where V refers to ex-post efficiency, co-pay decreases designed to raise adherence have opposing welfare effects on under and over-adhering patients as discussed.

Provider- or Manufacturer Reimbursement and Optimal Adherence

If adherence is driven by optimal learning, this has strong implications for the effects of various reimbursement policies set by payers to affect providers and manufacturers. Patients not adhering to poorly working therapies means that reimbursements are not spent on poorly performing care. This affects the impact of so called "pay-for-performance" schemes as well as explicit therapy stopping rules undertaken by providers or payers. In particular, stopping rules

imposed externally on patients only make them worse off in our framework. In addition, patient stopping rules mimic "risk-contracting" or pay-for-performance to manufacturers under which they only get paid when a therapy performs well on a population level. Patient learning implies such type of reimbursements may have small effects because payers do not pay for ineffective care when patients do not adhere to it.

Structural Estimation of Trial Attrition to Predict Future Adherence

The structural model of adherence discussed implies strong relationships between so called "real-world" vs clinical trial performance of treatments; sometimes distinguished by the names efficacy vs effectiveness. However, attrition behavior in clinical trials may stem from the same type of behavior as analyzed here (Philipson and DeSimone 1999). The central testable empirical implication of that past analysis as well as the adherence analysis here is that past performance drives current hazard rates into non-adherence. This prediction may be tested by longitudinal outcomes data in trials and data on both adherence and health outcomes in real world settings, the latter which may become more abundant as data on insurance claims and electronic medical records are merged.

Due to this similarity in behavior, there are conditions under which one can estimate the structural parameters from only having attrition behavior from clinical trial data. The parameters can then be used to predict or forecast post-marketing adherence behavior out-of-sample. In other words, structural estimation of attrition behavior in trials can allow for counter-factual predictions of future real world adherence and effectiveness.

In summary, we believe more explicit theoretical analysis of non-adherence would better expand our understanding of this important type of dynamic health care demand. Empirical

testing of explicit theories seems needed before making credible normative claims about the efficiency gains of various private or public interventions aimed at raising adherence.

References:

Baum, Stephanie. 2013. “Health IT startup claims pillbox app has boosted adherence rate to 81% in two months.” *MedCity News*, January 8, 2013. http://medcitynews.com/2013/01/health-it-startup-claims-pillbox-app-has-boosted-adherence-rate-to-81-in-two-months/

Bays, Harold E., Leiv Ose, Neil Fraser, Diane L. Tribble, Katherine Quinto, Robert Reyes, Amy O. Johnson-Levonas, Aditi Sapre, and Steven R. Donahue. 2004. "A Multicenter, Randomized, Double-Blind, Placebo-Controlled, Factorial Design Study to Evaluate the Lipid-Altering Efficacy and Safety Profile of the Ezetimibe/Simvastatin Tablet Compared with Ezetimibe and Simvastatin Monotherapy in Patients with Primary Hypercholesterolemia." *Clinical therapeutics* 26(11): 1758-1773.

Black, Dennis M., Richard J. Brand, Merwyn Greenlick, Glenn Hughes, and Jacqueline Smith. 1987. "Compliance to treatment for hypertension in elderly patients: the SHEP pilot study." *Journal of Gerontology* 42(5): 552-557.

Bosworth, Hayden B., Bradi B. Granger, Phil Mendys, Ralph Brindis, Rebecca Burkholder, Susan M. Czajkowski, Jodi G. Daniel et al. 2011. "Medication adherence: a call for action." *American Heart Journal* 162(3): 412-42.

Bradt, Russell N., S. M. Johnson, and Samuel Karlin. 1956. "On Sequential Designs for Maximizing the Sum of n Observations." *The Annals of Mathematical Statistics* 27(4): 1060-1074.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight, (2010), “Patient Cost-Sharing and Hospitalization Offsets in the Elderly”, *American Economic Review* 100(1): 193–213

Chernew, Michael E., Mayur R. Shah, Arnold Wegh, Stephen N. Rosenberg, Iver A. Juster, Allison B. Rosen, Michael C. Sokol, Kristina Yu-Isenberg, and A. Mark Fendrick. 2008 "Impact of decreasing copayments on medication adherence within a disease management environment." *Health Affairs* 27(1): 103-112.

Comstock, Jonah. 2013 “GlowCaps now sold through CVS, new randomized control trial launches.” *MobiHealthNews*, March 11, 2013. http://mobihealthnews.com/20750/glowcaps-now-sold-through-cvs-new-randomized-control-trial-launches/

Dickstein, Michael J. 2014. “Efficient Provision of Experience Goods: Evidence from Antidepressant Choice.” http://www.stanford.edu/~mjd/papers/eff_prov_experience _goods _mjdickstein.pdf

Feldman, Ross, Marilyn Bacher, Norman Campbell, Aidan Dover, and Arun Chockalingam. 1998. “Adherence to Pharmacologic Management of Hypertension.” *Canadian Journal of Public Health* 89(5): I16-18.

Flack, John M. , Serguei V. Novikov, and Carlos M. Ferrario. 1996. "Benefits of adherence to anti-hypertensive drug therapy." *European Heart Journal* 17(Suppl A): 16-20.

Gittins, John C., Kevin Glazebrook, and Richard Weber. 2011. *Multi-armed Bandit Allocation Indices, 2nd Edition.*, Wiley.

Gittins, John C. and David M. Jones. 1974. "A dynamic allocation index for the sequential design of experiments", in J. Gani, K. Sarkadi and I. Vince, eds. *Progress in Statistics,European Meeting of Statisticians 1972*, Vol 1 Amsterdam: North Holland, 241-266.

Gittins, John C. and David M. Jones. 1979. "A dynamic allocation index for the discounted multiarmed bandit problem." *Biometrika* 66(3): 561-565.

Goldman, Dana. Geoffrey F. Joyce, and Yuhui Zheng. 2007."Prescription Drug Cost Sharing" *Journal of the American Medical Association* 298(1): 61-69.

Goldman, Dana and Tomas J. Philipson. 2007. "Integrated Insurance Design in the Presence of Multiple Medical Technologies.", *American Economic Review* 97(2): 427-432.

GoodRx.com. 2014. "Simvastatin Prices." Accessed May 22, 2014, http://www.goodrx.com/simvastatin

Grossman, Michael. 1972. "On the Concept of Health Capital and the Demand for Health." *Journal of Political Economy* 80(2): 223-255

Haynes, R. Brian, K. Ann McKibbon, and Ronak Kanani. 1996. "Systematic review of randomised trials of interventions to assist patients to follow prescriptions for medications." *The Lancet* 348(9024): 383-386.

Heart Protection Study Collaborative Group. 2006. "Lifetime cost effectiveness of simvastatin in a range of risk groups derived from a randomised trial of 20536 people." *BMJ: British Medical Journal*, 333(7579): 1145-1148

Hershey, John C., Bruce G. Morton, Jane Braithwaite Davis, and Michael J. Reichgott. 1980 "Patient Compliance with Antihypertensive Medication." *American Journal of Public Health* 70(10): 1081-1089.

Hoadley, John F., Katie Merrell, Elizabeth Hargrave, and Laura Summer. 2012. "In Medicare Part D plans, low or zero copays and other features to encourage the use of generic statins work, could save billions." *Health Affairs,* 31(10): 2266-2275.

IMS Health. 2011. " The Use of Medicines in the United States: Review of 2010." Accessed May 14, 2014, http://www.imshealth.com/deployedfiles/imshealth/Global/Content /IMS%20Institute/Static%20File/IHII_UseOfMed_report.pdf

Jönsson, B., M. Johannesson, J. Kjekshus, A. G. Olsson, T. R. Pedersen, and H. Wedel. 1996. "Cost-effectiveness of cholesterol lowering Results from the Scandinavian Simvastatin Survival Study (4S)." *European Heart Journal,* 17 (7): 1001-1007.

Jovanovic, Boyan. 1979. "Firm-specific Capital and Turnover." *The Journal of Political Economy* 87(6): 1246-1260.

Lesselroth, Blake J., Patricia J. Holahan, Kathleen Adams, Zhen Z. Sullivan, Victoria L. Church, Susan Woods, Robert Felder, Shawn Adams, and David A. Dorr. 2011. "Primary care provider perceptions and use of a novel medication reconciliation technology." *Informatics in Primary Care* 19(2): 105-118.

Ljungqvist, Lars and Thomas J. Sargent. 2012. *Recursive Macroeconomic Theory, Third Edition*, MIT Press.

Mallion, Jean-Michel, Jean-Philippe Baguet, Jean-Philippe Siche, Frederic Tremel, and R. De Gaudemaris. 1998. "Compliance, electronic monitoring and antihypertensive drugs." *Journal of Hypertension. Supplement* 16(1): S75-79.

Moore, Michael J. and W. Kip Viscusi. 1988. "The Quantity Adjusted Value of Life." *Economic Inquiry*, 26(3): 369-388.

Moore, Michael J. and W. Kip Viscusi. 1990. "Models for Estimating Discount Rates for Long-Term Health Risks Using Labor Market Data." *Journal of Risk and Uncertainty*, 3: 381-401.

Murphy, Kevin M., and Robert H. Topel. 2006. "The Value of Health and Longevity." *Journal of Political Economy* 114(5): 871-904.

Nelson, Eugene C., William B. Stason, Raymond R. Neutra, and Harold S. Solomon. 1980 "Identification of the noncompliant hypertensive patient." *Preventive medicine* 9(4): 504-517.

New England Healthcare Institute. 2009. "Thinking Outside the Pillbox: A System-side Approach to Improving Patient Medication Adherence to Chronic Disease." Accessed June 2, 2014, http://www.nehi.net/writable/publication_files/file/pa_issue_brief_final.pdf

Pandey, Sandeep, Deepayan Chakrabarti, and Deepak Agarwal. 2007 "Multi-armed bandit problems with dependent arms." In *Proceedings of the 24th international conference on Machine learning*, (ACM): 721-728.

Pauly, Mark V. 1968. "The Economics of Moral Hazard: Comment." *American Economic Review* 58(3): pp. 531-37.

Pedersen, T., Kjekshus, J., Berg, K., & Haghfelt, T. 1994. "Randomised trial of cholesterol lowering in 4444 patients with coronary heart disease: the Scandinavian Simvastatin Survival Study (4S). *The Lancet*, 344: 1383-89.

Pharmaceutical Research and Manufacturers of America. 2013. "Current Organizations Funding Adherence Research.", accessed August 5, 2013, http://www.phrma.org/value/ sources-of-adherence-research-funding.

Philipson, Tomas J. 1997. "Data Markets and the Production of Surveys." *Review of Economic Studies* 64(1): 47-73.

Philipson, Tomas, and Jeffrey DeSimone. 1997. "Experiments and Subject Sampling."*Biometrika* 84(3): 619-630.

Philipson, Tomas J., and Larry V. Hedges. 1998. "Subject Evaluation in Social Experiments." *Econometrica* 66(2): 381-408.

Powell, Warren B. and Ilya O. Ryzhov. 2012. *Optimal Learning.* Wiley.

Rusmevichientong, Paat, and John N. Tsitsiklis. 2010. "Linearly parameterized bandits." *Mathematics of Operations Research* 35(2): 395-411.

Seabury, Seth A., Gupta, Charu, Philipson, Tomas J., Henkhaus, Laura E. and the PhRMA Medication Adherence Advisory Council. 2014 "Understanding and overcoming barriers to medication adherence: A review of research priorities," *Journal of Managed Care Pharmacy*, in press.

Smith, M.E. Beth, Nancy J. Lee, Elizabeth Haney, Susan Carson, Mark Helfand and Cathy Kelley. 2009. *Drug Class Review: HMG-CoA Reductase Inhibitors (Statins) and Fixed-dose Combination Products Containing a Statin: Final Report Update 5.* Oregon Health & Science University.

Viscusi, W. Kip and Michael J. Moore. 1989. "Rates of Time Preference and Valuations of the Duration of Life." *Journal of Public Economics,* 38: 279-317

Vollmer, William M., Adrianne Feldstein, David Smith, Joan Dubanoski, Amy Waterbury, Jennifer Schneider, Shelley Clark, and Cynthia Rand. 2011. "Use of Health Information Technology to Improve Medication Adherence." *The American Journal of Managed Care* 17(12 Spec No.): SP79-87.

Yao, Yi Ching. 2006. "Some results on the Gittins index for a normal reward process." in H. Ho, C. Ing & T. Lai, eds., *Time Series and Related Topics: In Memory of Ching-Zong Wei,* Institute of Mathematical Statistics, Beachwood, OH.

Yeaw, Jason, Joshua S. Benner, John G. Walt, Sergey Sian, and Daniel B. Smith. 2009. "Comparing Adherence and Persistence Across 6 Chronic Medication Classes." *Journal of Managed Care Pharmacy*, 15(9): 728-740

Zeckhauser, Richard. 1970. "Medical Insurance: A Case Study of the Tradeoff between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory* 2(1): 10-26