NON-ADHERENCE IN HEALTH CARE:

A POSITIVE AND NORMATIVE ANALYSIS

by

Mark Egan

and

Tomas J. Philipson[1]

The University of Chicago

January 9, 2015

[1] 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.

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 health care, totaling approximately 2.3% of GDP, and have proposed a plethora of interventions to raise adherence. However, a lack of explicit analysis exists in health economics of the important and dynamic demand behavior that drives non-adherence, and it is often attributed to uninformed or irrational patients. We argue that while providers may be more informed about the population-wide effects of treatments, patients are more informed about the individual value of a treatment. We interpret a patient's decision to adhere to a treatment regime as an optimal stopping problem where patients learn the value of a treatment through experiencing it. We derive strong positive and normative implications resulting from interpreting non-adherence as optimal stopping. Our positive analysis derives an "adherence survival function", depicting the share of patients still on treatment as a function of time, and predicts how various observable factors alter it. Our normative analysis derives the efficiency effects of non-adherence, the conditions under which adherence is too high as opposed to 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 even though only 43% of patients fully adhere.

Section 1: Introduction

A universally agreed upon and wide-spread challenge in health care is getting patients to better 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- or 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 that adherence is too low. 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 her individual value of a 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 approach of using a treatment; going on it to see seeing if it is valued by them, and terminating it if it is not. Indeed, non-adherence is 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 implying 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 also 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 sets up the model and derives the implied adherence survival functions. Section 3 discusses the large set of positive implications regarding the effects of observable factors on the adherence survival function. Section 4 discusses the normative implications for efficient adherence. Section 5 calibrates the size of

efficiency effects for the cholesterol reducing drug simvastatin. Lastly, Section 6 concludes with a discussion of several future research avenues suggested by the explicit analysis of this type of demand behavior.

Section 2: Non-Adherence as an Optimal Stopping Problem

In this section we derive the positive implications of interpreting non-adherence as an optimal stopping problem when a patient learns about her own personal treatment value given prior knowledge about population wide effects.[5] 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 depending on his personal valuation of 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.

We 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 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 health of the patient in a given period reflects the quality of the treatment plus some idiosyncratic shock (noise) according to

$$h_t = q + \varepsilon_t$$

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

Thus, the observable health of the patient depends on the unobservable quality of treatment q_t as well as other factors, ε_t, which makes inference about treatment quality and value of the treatment an issue. For example, the health of the patient may be determined by other factors than the quality 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 from the 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. 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. If not on treatment, the patient have access to an alternative treatment with per period utility $U(h_s, p_s)$. In the Appendix, we derive the analog implications for the case of multiple treatments which we interpret to include partial adherence levels of a single treatment.

Patients have a prior over the quality of the treatment, $F_0(\cdot)$. 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, perhaps through provider communication, 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, with many periods, 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 alternative treatment from there on. The future is discounted according the parameter β which may induce differential adherence across treatments with differences in timing of benefits and costs. 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.

Once a patient elects to forgo treatment for the alternative treatment, 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 alternative treatment.

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 (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 $S(1) - S(t)$.

---

[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.

Section 3: 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.

3.1 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 alternative treatment. The cost of treatment reduces the net benefit of treatment and ultimately lowers adherence-- 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 simply says that demand is downward sloping. 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. This has direct adherence implications for factors that affect prices.

The opportunity cost in terms of the alternative treatment, 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 alternative treatment raises adherence while the quality of alternative treatment 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 lowers adherence, and the relative price subsidies of coverage across treatments in a class affect adherence both in terms of the opportunity cost of learning but also through the value of learning.

3.2 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.

3.2.1 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. 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.

3.2.2 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.[8] These population-wide effects of treatments are often estimated in clinical trials conducted to

[8] Aggregate effectiveness should be distinguished from so called "cost-effectiveness" of care which is more loosely related to larger adherence. Cost-effectiveness is not perfectly related to adherence because patients may trade off cost versus effectiveness differently than one to one, that is, the utility $U(h, p)$ may differ from a 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.

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.[9]

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

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.

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

3.3.1 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

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

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 $\frac{dz_t|F_t(\cdot|\bar{h},\mu_0)}{d\sigma_\varepsilon^2} < 0$.[10]

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 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.

3.3.2 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 alternative treatment allows patients to partake in the upside of treatment quality without the downside risk. The option value of

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

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 $\frac{dz_t|F_t(\cdot|\bar{h},\mu_0)}{d\sigma^2} > 0$.[11] 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.

3.4: 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 the value of a treatment for a chronic condition such as e.g. diabetes.

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.

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

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

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).

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.[12] 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.

3.5: 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.

---

[12] 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.

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 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.[13]

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

[13] 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.

quality q rises. This therefore implies that real world effectiveness converges to trial effectiveness as the treatment improves

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

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

3.6: 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 $U(q, p) \geq U(q_s, p_s)$.

3.7: 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.

3.8: 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. It is important to stress that although a provider may communicate his expertise in the *population-wide effects* of the treatment, here F, the patient has the ultimate expertise in the *individual value* of the treatment after experiencing it. 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 the full costs of compliance and treatments. In particular, doctors may incorrectly 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. 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. For example, in the simple binary case the true value of the treatment is learned by the patient directly, so priors (whether influenced by providers or not) have no effect on secondary adherence.

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 may depend on the quality 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 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 4: Normative Implications for Efficient Adherence

In this section we discuss the efficiency implications of non-adherence.[14] Inefficient adherence transpires as the direct result of heterogeneous and unknown personalized treatment effects. The process in which individuals learn about their own value of the treatment creates the potential for both under and over-adherence.

---

[14] 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.

4.1 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 truly 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 alternative treatment

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

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

It is ex-post efficient for patients with personalized treatment effects above this reservation level to be on treatment; $q \geq q^*$. Therefore, there are two types of ex-post inefficiencies. The first inefficiency is under-adherence by those 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 those for whom treatment is not valuable. For the fraction

$F(q^*)$ of patients who do not respond, some will initially adhere to the treatment before learning that it is not valuable. [15]

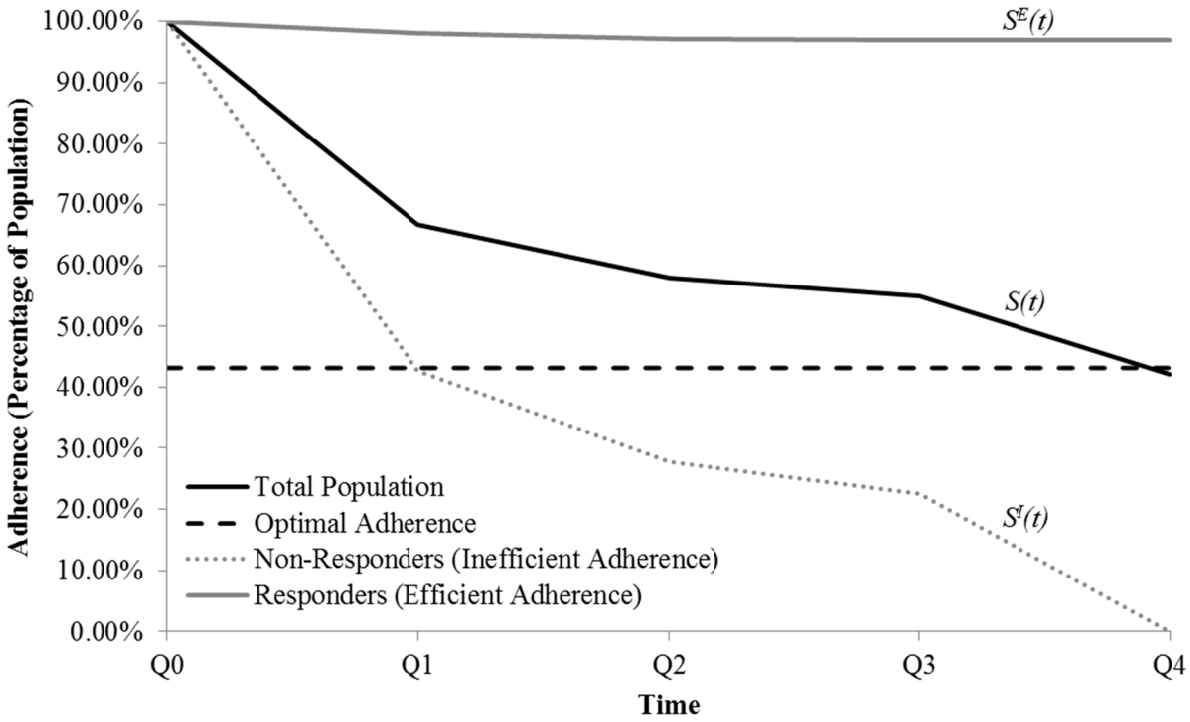
Figure 1 illustrates the general pattern of ex-post inefficient adherence. The sold gray line reflects the proportion of individuals at each period that adhere for whom the treatment is valuable and is reflected by the survival curve $S^E(t)$ indicating efficient adherence. It is the survival function for which the new treatment is preferred; $q > q^*$. 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 as reflected by the survival curve $S^I(\cdot)$ representing inefficient adherence. It is the survival function conditional for those with true quality $q < q^*$. Such inefficient over-adherence is reflected by a positive survival curve for this group that occurs because due to the option value of treatment or because 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; $\lim_{t\to\infty} S^I(t) = 0$.

---

[15] In this framework, traditional evaluations of adherence, that uniformly favors raising it, may be interpret to concern the special case involving the implicit assumption that the entire population benefit from treatment; q*=0. In This case, any factors that raise or lower adherence also raise or lower ex-post efficiency. However, when not everyone benefits from the treatment, factors that raise adherence may lower ex-post efficiency and factors that lowers adherence may raise such efficiency. The broken link between the level of adherence and ex-post efficiency is central to adherence interventions under patient learning.

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.

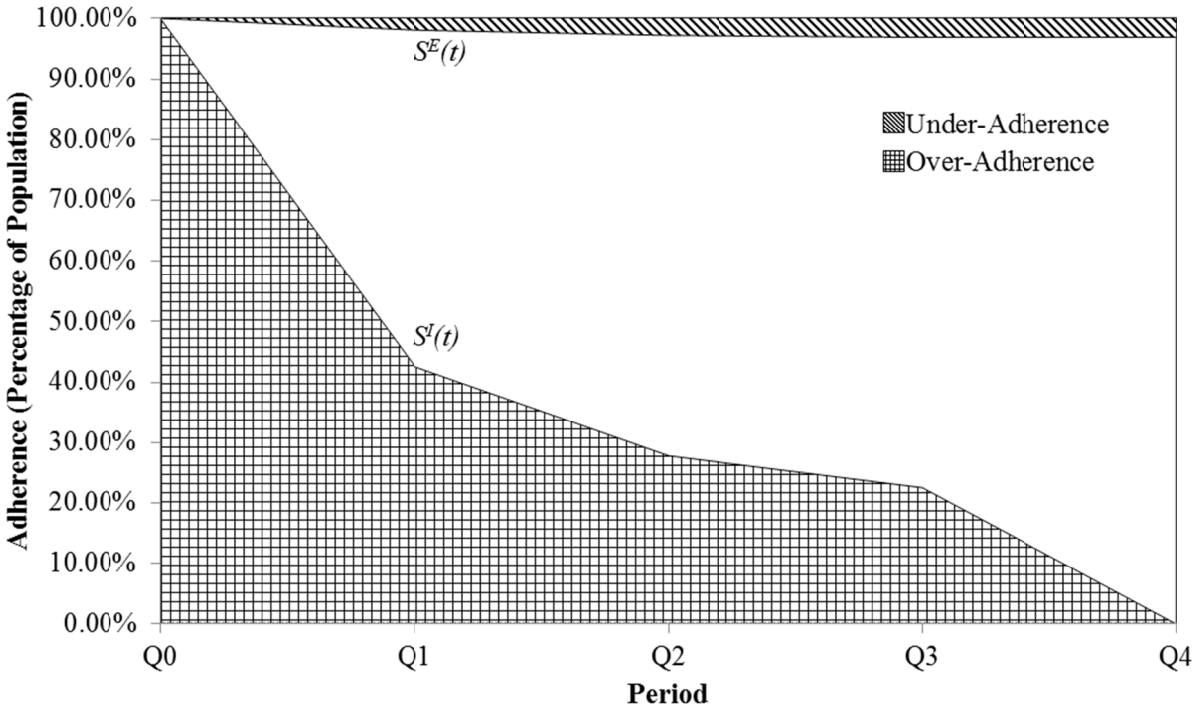
By our previous discussion that showed that true treatment quality raises adherence, the Figure depicts that efficient adherence is always larger than inefficient adherence; $S^I(t) < S^E(t)$, $\forall t$. The overall solid survival curve is the mixture of the two conditional survival functions with the mixture weights given by the fraction of true 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 (to match our discussion when we calibrate these curves later). It

follows that in the short run, there will initially be both under-adherence for true responders and over-adherence for true non-responders. 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 is decreasing and 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. 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 or integral 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 or integral below the non-responders survival function (S^I) represents over-adherence: fraction of patients that are adhering but should not be.

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 therapy is valued. Let the reservation price for the treatment for an individual of type q be denoted $p + r(q)$ and defined by

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

Thus the sign of $r(q)$ reflects whether the treatment is truly valued or not relative to the alternative treatment. It follows directly that higher performing treatments have a higher reservation price, $r'(q) > 0$, and that the price is negative for those who do not respond $r(q) \leq 0$ if $q \leq q^*$. The dollar value of the welfare loss at time t can then be written as

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

The first term $L_O(t)$ 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 $L_U(t)$ is the loss in welfare at time t from under-adherence; those who value the treatment but stopped adhering.

Given the welfare loss at each period, the present value of the total welfare loss over time is the discounted value of the loss from both forms of inefficiencies

$$L = \int_0^{\infty} \beta^t \left[L_O(t) + L_U(t)\right] dt$$

The important aspect of this overall welfare loss is that over-adherence is front-loaded while under-adherence that is back-loaded. This results from the fact that non-adherence occurs early during learning and vanishes in the long run. Therefore, in present value terms over-adherence often matters more than under-adherence. This is part of the reason we find that efficiency losses from over-adherence dominate under-adherence in our later calibration results for cholesterol lowering drugs. More generally, the importance of front-loaded over-adherence for overall efficiency goes against common arguments that there is under-adherence, though such arguments rarely or never have any explicit efficiency criteria stated or estimated.

4.2 Welfare Effects of Common Adherence Interventions

The analysis above implies that the efficiency effects of common adherence interventions are indeterminate even though such interventions may often raise overall adherence. We separate two types of interventions aimed at raising adherence; cost-based or education-based. Cost-based interventions aims to lower the time ort dollar expense of treatment while education-based interventions aims to raise the beliefs of the patient in the quality of the treatment or improve memory.

Consider a cost-based intervention represented by the parameter θ that raises survivals of all patients, at a cost $C(\theta)$. For example, this may be a reduction in copays or lowered travel costs to a provider. Raising adherence for everyone implies that over-adherence is increased and under-adherence decreased; $\frac{dL_0}{d\theta} > 0$ and $\frac{dL_U}{d\theta} < 0$. This implies counteracting effects on overall losses given by

$$\frac{dL}{d\theta} = \int_0^\infty \beta^t \left[\frac{dL_O(t)}{d\theta} + \frac{dL_U(t)}{d\theta} \right] dt$$

Therefore, any intervention that affects adherence behavior by the two groups *symmetrically* will have indeterminate effects on ex-post efficiency. On the other hand, this will likely be the norm as it is extremely difficult to make adherence interventions operate differently across the two groups. For example, a cut in co-pays raises efficiency for true responders, by making them adhere more when they should, but lowers efficiency for true non-responders, by making them also adhere more when they should not.[16] 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 calibrating these welfare effects in the proceeding section

The long-run effects of interventions raising adherence may be easier to sign. In the long run, we discussed over-adherence must go to zero when those truly not responding eventually learn that they should not adhere. Therefore, in the limit the welfare loss must fall from an intervention that raises adherence

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

Education and memory-based interventions are often used to attempt to raise adherence. However, unlike cost-based interventions they are hard to evaluate from an ex-ante efficiency perspective if the involve altering preferences through changed beliefs. Ex-post efficiency may

[16]More generally, 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[16] (Goldman and Philipson 2007).

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.

Section 5: Calibrating Adherence Inefficiencies: The Case of Simvastatin (Zocor)

In this section we calibrate our model of non-adherence to assess the welfare losses induced by ex-post inefficient adherence. We show how this is feasible given readily available data. We consider adherence associated with cholesterol lowering treatments taken by adult males. More specifically, we consider old males taking the drug simvastatin (Zocor) as a cholesterol lowering treatment regime. Our main result is that even though a majority of these patients do not adhere fully, the welfare loss of over-adherence dominates that of under-adherence. In particular, the loss due to over-adherence is over 80% larger than the loss due to under-adherence.

In our framework, we interpret simvastatin as the unknown treatment while the alternative is not taking any treatments. Our interpretation assumes that the sole objective of the treatment is to lower low-density lipoprotein cholesterol (LDL-C). The per-period (quarterly) benefit of simvastatin 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 $\varepsilon \sim N(0, \sigma_\varepsilon)$. Our calibration is for 58 year old males and we assume that patients expect to live the average life expectancy of 23 years without treatment but longer if responding to the simvastatin treatment.[17]

Patients observe their cholesterol levels through lab tests and update their adherence decision on a quarterly basis. We assume patients learn their true value of treatment fully 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.

A calibration of our model requires knowledge of the distribution of treatment effects, (μ, σ^2), how well the symptoms reflect treatment or signal noise (σ_ε^2), the costs of treatment (p), and the utility parameters (β, γ). Clinical trial data often provide information on the treatment quality and noise parameters. In particular, 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).

Table 1 summarizes the parameters values used in calibrating the model for simvastatin. We use clinical trial data on the distribution of effectiveness of simvastatin from (Bays et al. 2004). On average, the simvastatin treatment therapy in the Bays et al. study lowered LDL-C

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

levels by 37.00% over a quarter relative to no treatment (placebo).[18] In the context of our model 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) who using Medicare Part D data find that the median out of pocket cost paid by users for a one quarter supply of branded simvastatin was $231.25.[19,20]

TABLE 1: MODEL PARAMETER VALUES USED IN CALIBRATION

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. 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.

[18] Patients studied in Bays et al (2004) received dosages of 10-80 mg respectively of simvastatin per day.
[19] 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 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.

The health consumption trade-off parameter γ represents a patient's willingness to pay to lower his cholesterol for one quarter. The value of lowering cholesterol is induced from the value of the longevity increase it generates. In particular, is parameter is calculated using data on the longevity gains from simvastatin priced out to dollars using standard value-of-life estimates. Based on the results from the Scandinavian Simvastatin Survival Study (S4)[21], Jönsson et. al (1996) find that simvastatin treatment raised longevity by an estimated 0.377 undiscounted life years. These estimated longevity effects are in line with the results from the 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.[22,23] Standard existing estimates of the value of a life year at 81 ($VSLY_{81}$) is about $230,000 in 2014 (Murphy and Topel (2006)).[24] 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.

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

The value of lowering cholesterol can then be induced from the value of the longevity increase it generates. In particular, given our parameter estimates of the parameters μ, $VSLY_{81}$, and β

[21] 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.

[22] 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.

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

[24] 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].

above, the health consumption trade-off parameter satisfying equation (1) is $\gamma = 0.17\frac{\%}{\$}$. 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 simvastatin. Using claims data, Yeaw et al. (2009) implicitly estimate the adherence survival function we discussed for the cases of statins.[25] We calibrate β and $\sigma_0/\sigma_\varepsilon$ to minimize squared differences between the calibrated and empirical adherence survival function at each quarter for the first year.[26] The calibrated health discount factor is 0.90 which is line with the estimates from Moore and Viscusi (1988) and Viscusi and Moore (1989).[27] The calibrated signal to noise ratio is 2.00. This implies that an adhering patient's cholesterol level varies naturally from quarter-quarter with a standard deviation of $\sigma_\varepsilon = 7.40\%$.[28]

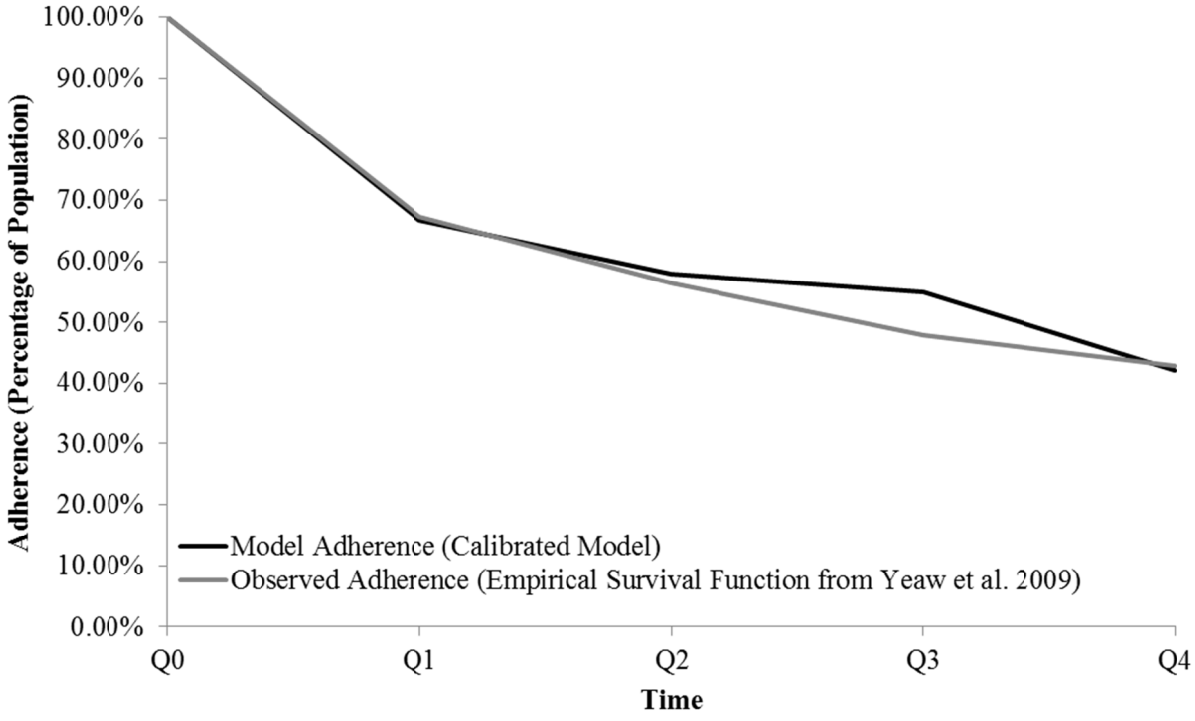
Figure 3 below display the adherence survival function from the calibrated model. It results from a simulation of 10 million hypothetical patients with the parameter values displayed in Table 1. The solid black line reflects the survival function corresponding to the calibrated model while the solid gray line reflects the observed adherence survival function as reported in

---

[25] 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).
[26] More precisely, we calibrate 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 $\sigma_\varepsilon/\sigma_0 \in \{0.50, 0.75, 1.00, 1.25, 1.50\}$.
[27] See Moore and Viscusi (1990) for further discussion on estimating discount rates for health outcomes.
[28] Note that in principle, the degree to which health symptoms reveal treatment quality (signal-to-noise 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 noise distribution, thereby allowing for out-of-sample predictions about future post-approval adherence behavior from trial data obtained pre-approval.

Yeaw et al. (2009). The calibrated survival function exhibits a correlation with the observed empirical survival function of 0.94.

FIGURE 3: CALIBRATED ADHERENCE SURVIVAL FUNCTION



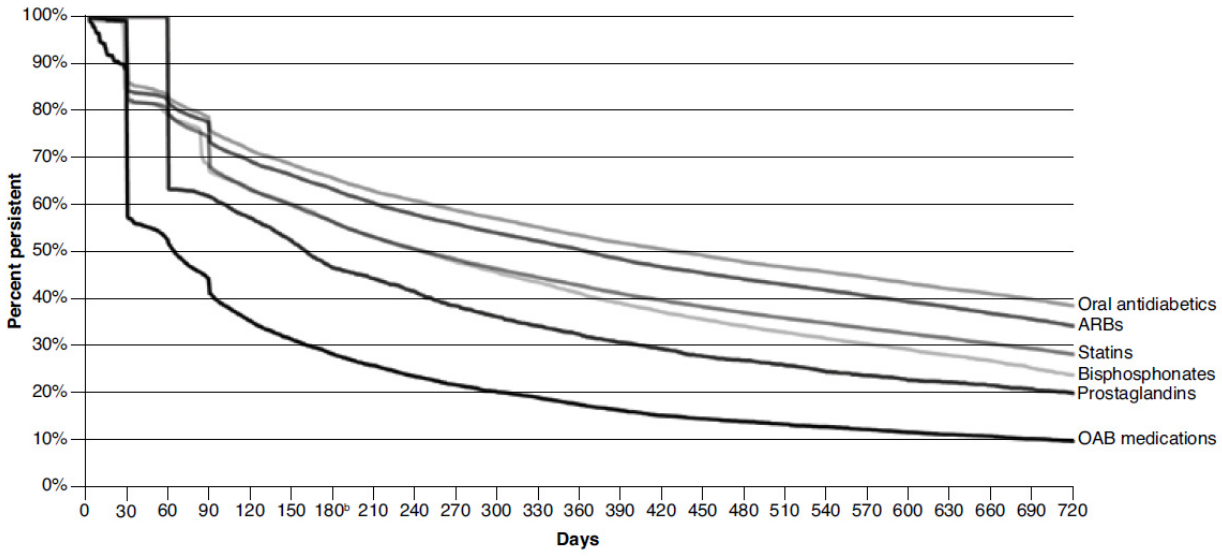
Notes: Figure 3 illustrates the calibrated survival curve corresponding to the parameter values in Table 1 and the empirical survival curve estimated in Yeaw et al. (2009).

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 in the long-run. The large degree of non-adherence is driven by the discounting of the fairly modest longevity effect of the treatment which is in addition discounted substantially because it occurs in the future. 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 these small 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. Overall, these simple facts suggest that low adherence to statins seem less puzzling than commonly discussed in both the economics and public health literatures.

The patterns of the observed adherence survival curve for simvastatin is displayed across many other treatments. Figure 4 below illustrates adherence curves for many alternative treatment categories.

FIGURE 4: ADHERENCE PATTERNS (YEAW ET. AL 2009)



Notes: Figure 4 illustrates observed adherence patterns for prostaglandin analogs, statins, bisphosphonates, oral antidiabetics, angiotension II receptor blockers (ARBs), and overactive bladder (OAB) medications. This figure is taken directly (Figure 2) from Yeaw et al. (2009).

The pattern that is common across these adherence survivals and that of simvastatin is that non-adherence is "front-loaded" in the sense that survival curves are convex and converge to a given

positive level. That is, non-adherence tends to disappear over time so that survival curves stabilize at a given level when hazard rates go to zero. These patterns are consistent with our discussed theoretical implications of more learning taking place early and that only true responders remain in the long run.

The calibrated model for simvastatin allows us quantify ex-post inefficient adherence behavior and separate the inefficiencies driven by over-adherence vs under-adherence. Indeed, Figure 2 in the previous section displayed the calibrated survival functions for non-responding and responding simvastatin users using the parameter values of Table 1. These two survival functions suggest that over-adherence may be more problematic than under-adherence. About 42% of true non-responders still take simvastatin after one quarter. However, as discussed, over-adherence vanishes as patients learn treatment as opposed to under-adherence that cannot be recovered. In the long run, only about 3.1% of true responders under-adhere to the treatment.

The calibrated model also allows us to dollarize the welfare losses associated with this inefficient adherence. Consider a patient adhering to the simvastatin treatment inefficiently which occurs whenever the health benefits do not exceed the cost of care; $\frac{1}{\gamma}q < p$. The associated welfare loss of over adherence is equal to the cost treatment minus dollarized health effect.

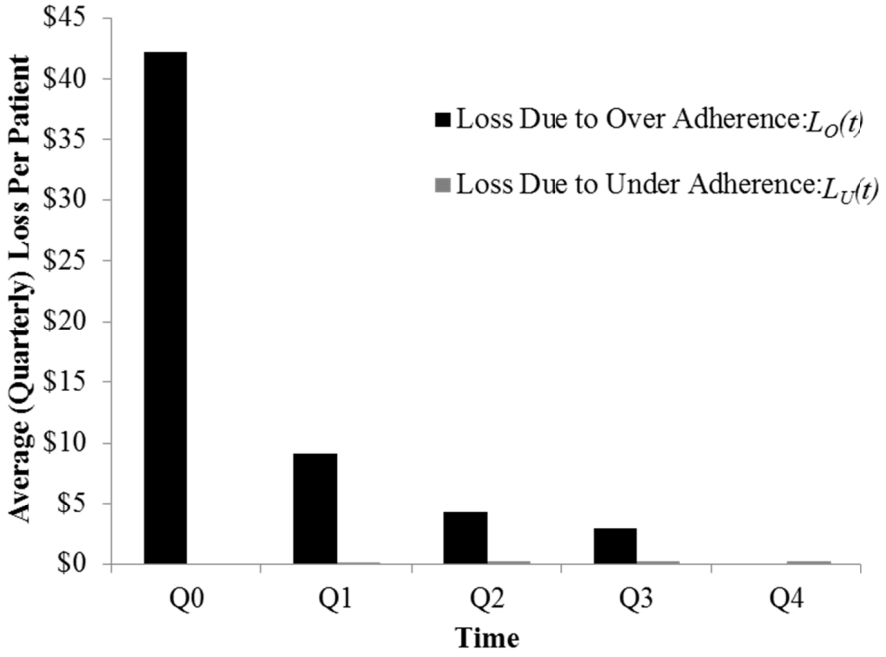
$$\text{Quarterly Loss Due to Over Adherence} = L_O = p - \frac{1}{\gamma}q$$

An analogous expression applies to those who truly value the new treatment but do not adhere

$$\text{Quarterly Loss Due to Under Adherence} = L_U = \frac{1}{\gamma}q - p$$

Figure 5 below displays the calibrated welfare losses from over- and under-adherence using these methods.

FIGURE 5: WELFARE LOSS OF EX-POST INEFFICIENT ADHERENCE TO SIMVASTATIN



Notes: Figure 5 illustrates the welfare loss due to over and under adherence corresponding to the parameters in Table 1. The loss due to 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 loss due to 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 5 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. As discussed, the loss from over adherence vanishes with time as patients learn that the treatment is not valuable. Since individuals never re-enter treatment, the loss from under-adherence rises over time but converges to a steady state level in perpetuity as those who value the treatment eventually stay on. The initial welfare loss due to over-adherence is $42.18 per patient-quarter but declines to zero as

non-responders drop out of treatment. Conversely, the loss due to 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 loss due to under-adherence is $6.52 per patient and for over-adherence is $57.87 per patient.[29] As discussed in the theoretical analysis, the larger present value effects of over-adherence steams from that it is front loaded in time as opposed under-adherence that is back-loaded. The total per capita loss due to inefficient adherence (over-adherence plus under-adherence) is $64.39. To put these numbers in perspective, a quantity of 94.1m simvastatin prescriptions were dispensed in the US in 2010 and the total spending on lipid regulators was $18.7bn (IMS Health 2011). The potential aggregate losses due to inefficient adherence are thus on the order of billions of dollars. Regardless, the major finding is that over-adherence losses greatly dominate under-adherence ones for statins, especially in present value terms.

Section 6: 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

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

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 analysis may usefully consider optimal insurance design when patient learning drives adherence. One direct way in which insurance design may affect adherence is through co-pays as generally recognized in the literature and as implied by the analysis above as well.

There is an interesting recent health policy literature on so called value based insurance design (VBID) which discusses copay design across treatments and services. 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 troublesome because 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.

This policy literature is in contrast to the 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 treatments and Goldman and Philipson 2007 for multiple ones). When value is defined as economic efficiency, then the VBID implied is that co-pays should be lower for less price-elastic treatments since the insurance it enables is less counteracted by inefficient demand at below marginal-cost prices. Indeed, the VBID induced by the traditional economic analysis implies that if one were to marginally change co-pays to alter adherence it would result in inefficient adherence. It is not clear why suggested benefit re-designs of the more recent policy literature on VBID, often arguing that copays should be raised or lowered to alter adherence, affects efficiency favorably nor any other definition of V.

A better understanding of how patient learning affects VBID is a useful area of future research. It is generally true that for experience goods such as medical treatments, initially low prices may subsidize learning efficiently. Our analysis has several implications for optimal co-pay designs that may be usefully addressed in future research. One possibility is the value of temporary copay rebates to prevent under-adherence by those who benefit from care. The role of pharmaceutical samples being provided free of charge in mitigating inefficient adherence is of particular importance to better understand.

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

Appendix

Multiple Treatments and Partial Adherence on a Single Treatment

This Appendix generalizes the discussion to the case when the quality of more than one treatment is unknown to the patient. Instead of two treatment alternatives, one uncertain treatment and one certain alternative treatment, the more general setting with K uncertain treatment alternatives is considered. One can think of the K treatments as completely different treatments or alternatively as different levels of adherence with the same treatment.

Multiple uncertain treatments involve a so called multi-armed bandit problem in statistical decision theory. 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.[30] 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

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

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

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

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 a single uncertain treatment ($k = 1$) and a certain alternative treatment ($k = 2$) corresponds to K=2 with $\sigma_2^2 = 0$.

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 alternative treatment for which the patient is indifferent between treatment and the alternative treatment in the simple two treatment alternative case.[31]

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

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

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

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.

Implications for Non-Adherence among Multiple Treatments

Computing Gittins indices correspond directly to the single armed bandit problem and optimal stopping rules described for the two treatments 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), the comparative statics discussed 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 ($\sigma_{k\varepsilon}^2$).

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.