Innovation and Diffusion of Medical Treatment*

Barton H. Hamilton
Olin Business School, Washington University in St. Louis

Andrés Hincapié
Department of Economics, University of North Carolina at Chapel Hill

Robert A. Miller
Tepper School of Business, Carnegie Mellon University

Nicholas W. Papageorge
Department of Economics, Johns Hopkins University and IZA

This version: January 8, 2018

Abstract: We develop and estimate a dynamic structural model of demand for a product line whose characteristics evolve over time as a consequence of consumer choices. We provide a new approach to the econometric challenge of estimating demand under uncertain innovation that includes sporadic breakthroughs and frequent, incremental changes. We use our framework to analyze consumer choice and the realized path of innovations over a long time horizon in a maturing product market: HIV drugs. In our model product quality is multidimensional since medications differ by their efficacy and their propensity to cause side effects. We allow for the possibility that new, more effective medicines can sometimes have harsher side effects. Atomistic consumers do not account for the role of aggregate demand on the speed and direction of innovation, leading to possible externalities. Using our estimated model we find that a planner that internalizes the externalities can increase welfare by at least 2% by increasing experimentation. Our results also indicate that providing monetary incentives for trial participation can be welfare improving.

Keywords: Innovation, Dynamic Demand, Structural Models, HIV/AIDS, Clinical Trials.
JEL Classification: O31.

*This paper has benefited from presentations in the economics departments at Arizona State University, University of Toronto, University of Central Florida, University of Washington, New York University, University of Wisconsin - Madison, Yale University, Monash University, University of Sydney and in business schools at Washington University in St. Louis, Carnegie Mellon University, Stanford, and Harvard, as well as a conference presentation at the 2016 Annual Health Econometrics Workshop. We also thank John Rust for his comments.

1 Introduction

In many product markets, innovation leads to substantial changes in product quality from one point in time to the next. Research relating innovation to consumer demand tends to emphasize the impact of new products on individuals' choices, behavior and welfare. Going back to Hicks (1932), economists have recognized that consumer behavior can also drive innovation. The idea is that innovation is not a random scientific process. Rather, profit-seeking firms direct their inventive activity (e.g., R&D investments) in an effort to develop new products that meet potential demand. This type of demand-driven innovation (sometimes referred to as "demand pull" (Schmookler, 1966; Scherer, 1982)) implies a possible externality since atomistic consumers do not account for the impact of aggregate behavior on the innovation process when making decisions (Jovanovic and MacDonald, 1994; Waldfogel, 2003; Finkelstein, 2004; Goettler and Gordon, 2011). A potential implication is that pricing the externality could improve consumer welfare by accelerating the introduction of new products.

Several features of the market for pharmaceuticals make it an interesting context to study demand-driven innovation. First, medical products have two dimensions of quality: efficacy and side effects and consumer preferences are heterogeneous. Thus, it is not generally meaningful to see one product as strictly better than another and many differentiated products can coexist in a given market. Innovations can also be better on one dimension and worse along another, a leading example being effective new medicines with harsh side effects. Second, in medical markets product quality is often uncertain, especially when products are new. Experimentation is therefore common among consumers and helps to drive innovation (Bolton and Harris, 1999; Dranove et al., 2014). Patients often resort to experimenting when they are sick and lack access to better options. It is thus often the most desperate among consumers who drive medical innovation, which benefits healthier patients along with generations of potential future patients.[1] Correctly pricing the externality could thus not only improve efficiency, but also equity. The reasoning is that incentivizing experimentation among all consumers (rather than relying solely upon the sickest patients) could accelerate innovation as in other markets, but also distributes the burden of innovating more evenly across patients.

In this paper, we specify and estimate a structural model of demand-driven innovation. In the model, individual consumers maximize their lifetime utility by choosing from a menu of medical treatments, which includes the option to experiment with new products. Consumers

[1]This point is linked to the model of endogenous growth in Romer (1986) where producer innovations may generate profits for potential future producers.

experiment by choosing products they have never tried before and become fully aware of their qualities only after they have used a product at least once. Alternatively, they can use new technologies that are not yet on the market. In software, this is known as beta-testing; in medicine, this is done through participation in clinical trials. Experimental products may be superior to products already on the market, but they may also be of dangerously low quality.

A key feature of the model is that innovation is endogenous to aggregate demand. In particular, we estimate a stochastic process of product innovation that is conditioned on shifting market shares, including rates of trial participation. Since products are multi-dimensional, market shares affect can affect both the speed of innovation and also its direction, i.e., the relative magnitude of shifts in efficacy versus side effects given the current state of medical technology. When making choices, consumers use this innovation process to form beliefs about future innovations and resulting evolving choice sets. Thus, consumers are neither fully aware of how the product market will evolve, nor fully unaware, in which case technology shifts would be treated as unexpected regime changes. Consumers also account for the impact of aggregate behavior on future products when making choices. To fix ideas, consumers in relatively good health may face incentives to avoid new or experimental medical treatments if they expect large numbers of other consumers to participate in clinical trials and thus generate better drugs in the future.

We match our model to data on the realized path of innovations, product quality and consumer choices over a long time horizon in a maturing product market: HIV drugs.[2] HIV is a medical condition that reduces the ability of the immune system to fight off routine infections (a condition known as AIDS).[3] It reached epidemic proportions in several countries starting in 1984 leading to just over 613,000 deaths in the U.S. by 2008.[4] The benefit of observing a long panel in the market for HIV drugs is that we can see how the path of innovation unfolded over time. In developed countries, where access to medication is widespread and subsidized, technological advancement means that HIV is currently a manageable condition and the side effects of medications are fairly mild. This was not always the case. In the early years of the epidemic, available treatments were not only largely ineffective, but also had uncomfortable, painful and even deadly side effects. Each year brought innovations, most of which were small. Some new medications were worse than existing technology since

[2]HIV stands for human immunodeficiency virus.

[3]AIDS stands for acquired immunodeficiency syndrome.

[4]For comparison, over the same period in the U.S., there were 508,000 homicides and U.S. deaths in World War II were just under 420,000. Currently, there are roughly 50,000 new infections and 13,000 deaths per year in the U.S. that are attributed to HIV/AIDS. Globally, the number of deaths due to HIV/AIDS stands at roughly 35,000,000.

since they were more toxic without being more effective. In the mid-nineties, a new set of treatments (collectively known as HAART) was introduced, which transformed HIV from a virtual death sentence to a chronic condition.[5] Within two years, mortality rates fell by over 80% among HIV+ men (Bhaskaran et al., 2008). HAART therefore marked a clear departure from existing products in the market for HIV treatments. However, HAART also involved drugs that were highly toxic, leading to side effects that were often intolerable and drove some people to avoid using them. Strikingly, innovations occurring after HAART had fewer side effects, but were generally not more effective than earlier versions of HAART. To explain this pattern, we examine the possibility that, once a level of efficacy was achieved that insured patient survival, patient demand for drugs with fewer side effects helped to drive innovation towards treatments with fewer side effects.

Our results reveal that individuals' preferences tilt the path of innovation towards treatments with fewer side effects, away from the invention of more effective treatments. Moreover, individuals have a strong distaste for experimentation that can slow the diffusion of new, superior products as well as the development of future treatments in clinical trials. Because individuals are atomistic and they do not incorporate the consequences of their actions in the path of technology, a social planner could improve welfare by fostering experimentation. As a measure of the externality, we compute the marginal increase in aggregate welfare generated by a planner who sends the marginal person to clinical trials at the atomistic equilibrium. Since the marginal person does not want to join a trial, he loses a little more than \$600 when he is forced to participate. However, because trial participation spurs innovation by pushing up the expected quality and the expected number of new products, the net social gain is about \$2,000 per individual. Finally, we find that a more realistic policy that relies on a flat subsidy to induce a marginal increase in trial participation also generates net social gains. Therefore, our results indicate that providing monetary incentives for trial participation can be welfare improving by accelerating the progress of innovation.

This study contributes to a literature on dynamic demand under uncertainty. Following Petrin (2002), each product in our model is a bundle of characteristics, in our case, efficacy and side effects.[6] Moreover, similar to (Gowrisankaran and Rysman, 2012), we allow product characteristics to have dynamic impacts on consumers. An important feature of our model is that we allow rational consumers to experiment with new products. In our context, there

[5]HAART stands for highly active anti-retroviral treatment. There is no vaccine or cure for HIV or AIDS, but HAART is the current standard treatment. In general, 1996 is marked as the year when two crucial clinical guidelines that comprise HAART came to be commonly acknowledged. First, protease inhibitors (made widely available towards the end of 1995) would be an effective HIV treatment. Second, several anti-retroviral drugs taken simultaneously could indefinitely delay the onset of AIDS.

[6]Studies pioneering the 'characteristics approach' include Stigler (1945), Lancaster (1966) and Rosen (1974).

are two motivations. First, following Erdem and Keane (1996) and Crawford and Shum (2005), we allow consumers to gain by experimenting with unfamiliar products, which they can continue to use in the future.[7] Second, we allow dynamically optimizing consumers to participate in clinical trials to use products that are otherwise unavailable, may be worse than the state of the art, but may provide early access to lifesaving new technologies. Modeling trial participation as a rational choice relates our work to Chan and Hamilton (2006), who model the decision to remain in a clinical to maintain access to good HIV medicine.

A key departure from literature on dynamic demand with experimentation is that we explicitly model how these decisions driven innovation and thus future products, which consumers forecast when making their current decisions. Several papers have demonstrated that market size affects the speed of innovation. For example, Finkelstein (2004) shows that policies promoting vaccine use accelerate the development of vaccines. Also in the medical context, Dranove et al. (2014) identify a "social value" of pharmaceutical innovation, showing that Medicare Part D spurred the development of some drugs. A common idea in this literature is that if consumer behavior drives innovation, which benefits other consumers, it follows that a demand externality arises. Waldfogel (2003) uses the term "preference externalities" to describe the mechanism through which market shares can influence products, thus benefitting consumers with similar tastes.[8] More closely related to us, Bolton and Harris (1999) argue that a free-riding problem emerges if experimentation accelerates innovation. In our context, if clinical trials provide social benefits by spurring innovation, individually rational consumers may choose to participate less than is socially optimal.

We also contribute to research on structural estimation by providing a simulation-based econometric method to estimate models of endogenous innovation. Methodologically, our approach builds on Hotz and Miller (1993) and Hotz et al. (1994) in using conditional choice probabilities (henceforth, *CCPs*) and forward simulation techniques to incorporate how individuals form expectations about future innovations.[9] In our context, the choice set that individuals face evolves stochastically over time, which means the problem is non-stationary. To make our model tractable, we summarize the current state of technology using a non-stationary reference point (or *centroid*) that emerges endogenously from consumer

[7]Empirical models of learning and experimentation also include Miller (1984) and Hincapié (2017).

[8]He also highlights the individuals with different tastes benefit less. Demand externalities have been discussed in a variety of scenarios, including sorting into neighborhoods (Bayer and McMillan, 2012) and the emergence of food deserts (Allcott et al., 2015). In the context of obesity, Bhattacharya and Packalen (2012) provide evidence that individual efforts to prevent obesity can shrink the market size for obesity treatments, which slows technological progress. If so, individuals may over-invest in preventative care compared to the social optimum.

[9]We also build on Altuğ and Miller (1998) in providing an empirical dynamic model with aggregate shocks.

demand. We then define a non-parametric, stationary distribution of innovations identified from the distance between the centroid and new products in the following period. We model consumer beliefs about future choice as technology paths simulated using this innovation distribution.

Our approach to handling non-stationarity in modeling innovation is similar to the approach in Goettler and Gordon (2011), who develop a model relating market structure to innovation in the market for microprocessors. They find that, in contrast to a monopoly, the presence of a second firm can slow innovation (since firms do not expect to capture all profits), but that consumer surplus falls in the absence of a competing firm due to monopolistic prices. To handle non-stationarity, in each period, the state of the art becomes the starting point for future innovations, which is analogous to the the role the centroid plays in our model. Despite these similarities, there are some important differences in their setting and thus their modeling choices. In their setting, product quality is one-dimensional and the innovation distribution is effectively binary (either improving by a fixed amount or not). They also assume consumers are homogeneous, which means that the choice set in their context is also effectively limited to upgrading to the best technology or staying with the current one. In contrast, in our case, product quality is multi-dimensional, which means that product quality can evolve in many different directions on a two-dimensional plane. Also, as we show, the empirical distribution of innovations for HIV drugs is not well-approximated as movements with a fixed distance. Finally, we must account for a larger choice set since multiple dimensions of product quality coupled with consumer preference heterogeneity imply that many products can co-exist in a single market. A weakness in our approach in comparison to Goettler and Gordon (2011) is that we are unable to explicitly model firm interaction, which limits counterfactuals we can perform. A benefit of our approach, however, is that we are able to examine welfare implications of counterfactual policies in a setting where consumer behavior not only affects the speed of innovation, but also the direction it takes by tilting the path of technology towards more favorable products, in our case, those with lower efficacy and fewer side effects.[10]

The remainder of this paper is organized as follows. Section 2 describes the data set we use. In Section 3, we specify the structural model and in Section 4 we discuss estimation. In Section 5, we present parameter estimates and describe model implications for the distribution of innovations. In Section 6, we study counterfactual technology paths and the link between consumer choice and innovation. In Section 7, we examine the choice externality

[10]In particular, we are unable to conduct policy analysis related to market structure using our framework. An interesting extension of the current paper would be to merge the two approaches by integrating firm decision-making into a model where products have multiple qualities.

and consumer welfare. Section 8 concludes.

2 Data

In this section we introduce the data set used in this paper and describe some of the key empirical patterns we use to identify structural parameters. We use the public data set from the Multi-Center AIDS cohort Study (MACS). The MACS is an ongoing longitudinal investigation (beginning in 1984) of HIV infection in men who have sex with men (MSM) conducted at four sites: Baltimore, Chicago, Pittsburgh and Los Angeles.[11] At each semi-annual visit, survey data are collected on HIV+ men's treatment decisions, out-of-pocket treatment expenditures, physical ailments, which can reflect drug side effects, along with sociodemographic information, such as labor supply, income, race, and education.

In addition, blood tests are administered at each visit to objectively measure health status. Our main objective measure of immune system health is *CD4 count*, defined as the number of white blood cells per cubic millimeter of blood. Absent HIV infection, a normal range is between 500 and 1500. For HIV+ individuals, a count below 500 indicates that the immune system has begun to deteriorate due to HIV, but can still fight off infections such that the individual is not symptomatic. When CD4 count drops below about 300, a patient is said to suffer from AIDS.[12] AIDS means that the immune system becomes unable to fight off routine infections and survival probability drops. The MACS data set is particularly well-suited for an analysis of demand-driven innovation. Few data sets have a continuous, precise measure of underlying health, additional data on physical health outcomes, detailed treatment data along with information on economic outcomes.

2.1 Summary Statistics

The full MACS data set contains information on 6,972 subjects at 49 possible semi-annual visits for a total of 111,271 observations in the form of subject-visit dyads. We limit our

[11] Data in this manuscript were collected by the Multi-Center AIDS Cohort Study (MACS) with centers (Principal Investigators) at The Johns Hopkins Bloomberg School of Public Health (Joseph B. Margolick, Lisa P. Jacobson), Howard Brown Health Center, Feinberg School of Medicine, Northwestern University, and Cook County Bureau of Health Services (John P. Phair, Steven M. Wolinsky), University of California, Los Angeles (Roger Detels), and University of Pittsburgh (Charles R. Rinaldo). The MACS is funded by the National Institute of Allergy and Infectious Diseases, with additional supplemental funding from the National Cancer Institute. UO1-AI-35042, 5-MO1-RR-00052 (GCRC), UO1-AI-35043, UO1-AI-35039, UO1-AI-35040, UO1-AI-35041. Website located at http://www.statepi.jhsph.edu/macs/macs.html.

[12] AIDS stands for acquired immunodeficiency syndrome. The CD4 cutoff below which AIDS occurs varies between 200 and 350.

attention to HIV+ individuals, leaving us with 47,753 observations. Due to lack of data on gross income and out-of-pocket treatment costs at earlier visits, we drop observations prior to visit 14 (roughly, late 1990) and for robustness in the reporting of survival we also drop observations after visit 47 (about 2008). These sample period restrictions leave us with 29,523 observations and 2,420 individuals. Next, we drop observations where data are missing on at least one of the variables used in subsequent analysis (though we conduct various robustness checks to insure that our results are not driven by these exclusions). After these exclusions, the remaining analytic sample consists of 1,719 unique individuals and 16,851 observations.

Summary statistics by individual are reported in Table 1. The first column presents statistics for the analytic sample.[13] 68% of sample subjects are white, 22% are black and about 9% are hispanic. Race variation in our sample is important since previous research has emphasized difficulties in recruiting blacks into clinical trials, which may reflect different costs associated with treatments or variation in expected health outcomes (Harris et al., 1996). About 86% of the sample received some secondary education or more and nearly a quarter (23%) attended graduate school. Consistent with previous research studying medication choice using the MACS data set, there is evidence of substantial variation in labor supply (Papageorge, 2016). 74% of the sample is observed working at least once and 68% of the sample is observed not working at least once.

Underscoring the seriousness of HIV infection, about 40% of the HIV+ subjects we observe at least once over the sample period die prior to the end of the sample period. However, product market innovation led to drastic changes for HIV+ men. The most striking example is the introduction of HAART in the mid-1990s, which was much more effective at improving underlying health compared to the treatments that preceded it. Conditional on surviving until the invention of HAART, 20% of subjects are observed dying. This understates the impact of HAART since the sample under study is an aging cohort, i.e., observed survival rates are much higher even when the cohort is older after HAART becomes available. Further, according to Table 1, about 83% of subjects are observed using a market product at least once. Moreover, nearly a quarter (24%) opt for early access by participating in a clinical trial at least once during the sample period, suggesting that patients are willing to try experimental products where quality is uncertain.

[13]For comparison, the third column reports statistics for a larger sample of 2,420 individuals, where we have not dropped observations due to missing data on any particular variable.

2.2 Consumer Demand

In this section, we study consumer demand in the maturing market for HIV drugs. We emphasize two key patterns in the data. First, consumers are willing to use drugs with side effects when drugs are also effective. Otherwise, they often avoid drugs altogether. Second, consumers participate in clinical trials when they are very sick and when existing technologies are of low quality. Once good technology comes available, willingness to experiment plunges. Together, these patterns in the data support two ideas that underlie our theoretical model developed in Section 3. First, product quality in the medical context is multi-dimensional. Second, experimentation is a rational choice to gain access to unavailable and possibly superior technology.

In conducting our preliminary analysis of consumer demand, we pay close attention to comparisons of behavior before and after the introduction of HAART. Since HAART marked a large innovation on earlier treatments, it induced strong and observable consumer responses that help to identify consumer preferences over medications. Summary statistics for subject-visit dyads are found in Table 2 for the full analytic sample (column [1]) and then separately for the pre and the post-HAART eras (columns [2] and [3], respectively). We split the sample by HAART era to illustrate substantial changes to choices and outcomes after HAART was introduced.

Perhaps the most striking example of the impact of HAART on consumers is through its effect on survival. In Figure 1, we plot the probability of dying between periods t and $t+1$ conditional on survival until t. Death rates are much higher prior to HAART introduction and despite a multitude of new treatments coming available. After HAART, death rates plunge, and continue to fall until 2007, as smaller innovations occur that make drugs incrementally more effective and less toxic. HAART introduction also affected immune system health, as measured by CD4 count. According to Table 2, average CD4 count among HIV+ men in our sample is 407 in the pre-HAART era, rising to 524 in the post-HAART era. In Figure 2(a), we plot average CD4 count over time for people on market drugs and no treatment for HIV. Over time, health for people taking no drugs remains fairly constant while health for individuals in a market drug rises.[14]

Given the impact of HAART on health, it is important to understand why many consumers did not use it. In Figure 3(a), we plot the proportion of HIV+ consumers using an HIV treatment. Notice that treatment consumption is about 50% in 1990 and actually falls prior to HAART introduction. This reflects that products available on the market are

[14]Notice that average age rises and labor supply and income decline after HAART, consistent with the fact that we observe an aging cohort, which is more likely to retire and report lower gross income over time.

of fairly low quality. Still, if quality were uni-dimensional, even a low quality drug would be better than no drug at all. Moreover, even after HAART is invented, though there is a considerable rise in market product usage, there is a substantial proportion of HIV+ men not using treatment.

Treatment costs are one possible explanation. In Table 2, we see that treatment costs rise after HAART introduction, from about $179 to $327 for six months of treatment. In other words, even in the post-HAART era, costs are fairly low given that individual earnings average about $37,000 per year. It is worth mentioning, moreover, that non-users of market drugs pay non-zero costs for drugs, perhaps spending more money on medication to fight opportunistic infections. In other words, the incremental out-of-pocket cost of effective HIV treatments does not appear sufficient to explain why some people avoid HIV treatments.

Another possibility is that drug quality is multi-dimensional in which case demand reflects a distaste for another feature of HIV drugs. Given data on physical ailments, we explore the possibility that consumer demand reveals a distaste for side effects. Interestingly, after HAART introduction, the proportion of individuals reporting physical ailments declines only slightly (45% to 41%). The small change reflects the net effect of two countervailing dynamics (Papageorge, 2016). HAART improved health on average, which lowered reported ailments attributable to symptoms of HIV. However, HAART also led to side effects among users, thereby increasing reports of ailments. The increase in side effects also reflects how use of HIV treatment rose with the introduction of HAART, from 45% to 76%. We also plot physical ailments over time in Figure 2(b). For non-users of HIV medications, ailments remain fairly steady. For users of HIV medications, ailments drop prior to HAART introduction and then rise after HAART, which is consistent with HAART being a highly effective drug with side effects. However, after 2001, ailments decline for individuals using HIV drugs. This reflects later improvements to medications, which lowered their side effects.

Further evidence in support of the idea that there are two important dimensions of quality that influence demand comes from market consumption by CD4 count, plotted in Figure 3(a). Sicker people are far more willing to take low effective medications despite side effects in the years before HAART. After HAART, notice a striking convergence in the proportion of men using medications, driven largely by healthy individuals going onto medication. Thus, the rise in consumption of HIV treatments after HAART was introduced suggests that patients are more likely to use drugs despite side effects if the utility cost of suffering ailments is offset by expected improvements to health. HAART was more effective than earlier drugs, which encouraged people to use it despite its side effects. This would explain the rapid rise in use of HIV treatments after HAART is introduced since individuals would be more willing to use drugs with side effects as long as drugs are effective at improving underlying health.

Another option for individuals in the product market we study is to join a clinical trial to gain early access to new products. Studying how individuals experiment with new drugs by joining a clinical trial further highlights how consumers respond to innovations in the market for HIV drugs. Trial participation over time and by health status is plotted in Figure 3(b). The figure reveals several dynamics. First, early trial participation is driven largely by individuals with low CD4 counts. This suggests that, as individuals become ill, they also become more willing to experiment with new products of uncertain qualities. Second, in the years just prior to HAART introduction, the drugs that comprise HAART, including protease inhibitors, marked a substantial improvement over drugs available on the market. In those years, trial participation gave individuals early access to much better products. This relates to the idea of *beta testing* in markets where some consumers are willing to experiment with new products with high potential quality.

After HAART trial participation plunges and there is a marked convergence by health status in the proportion of patients in trials. These patterns suggest that, once effective drugs are available, trial participation is no longer driven by sick people willing to face uncertainty in exchange for early access to a possibly high-quality product. The reason is that, once HAART is available on the market, individuals no longer need to participate in a trial to access good drugs. Shifts in trial participation over time suggest that this form of experimentation is a rational choice to gain access to new technology, especially in a maturing market where existing technology is not particularly good. Patients, especially sick ones whose survival is at risk, are willing to experiment to gain access to something better, but are less willing to do so when good treatment is available on the market.

2.3 Market-Level Innovation

The previous section examined consumer demand patterns in light of shifting product quality over time, highlighting that side effects seem to play a larger role in demand after survival is more or less assured. In this section, we consider market-level innovation. To start, we illustrate innovation and diffusion of new products over time in the market for HIV treatment using a "heat map" displayed in Figure 4. For the approximately 90 drugs that were most used, we compute market share over our sample period.[15] Dark blue corresponds to low (or zero) market share and warmer colors indicate higher market shares. Several patterns

[15]Appendix A is a data appendix that contains additional information on individual drugs and treatment combinations. Table **??** discusses which drugs or combinations are taken in clinical trials. Table S1 lists the chemical compositions of each drug. Table S2 shows how some drugs are combined into treatments. Table S3 presents our market products, which are the main sets of treatments we observe, including the individual drugs they are composed of, whether or not they count as HAART and their entry and exit visits.

emerge from this heat map. In earlier years, there are fewer treatments with high market shares. Over time, as many treatments are introduced, market shares drop, which suggests there is heterogeneity in preferences. In fact, low market shares are common in the years following HAART introduction, when many new treatments were introduced, most of which were effective, but with strong side effects. After HAART, moreover, many drugs became obsolete, suggesting that new drugs are improvements on old ones.[16] As the market matured, some treatments were developed that were effective and offered fewer side effects, yielding a concentrated market once again.

2.4 Relating Innovation to Demand

Finally, we discuss evidence for the idea that the observed innovation path is a response to consumer preferences. In Figure 5, we plot drug qualities (effectiveness and side effects) for different periods of time. The figure illustrates the path of technology over time. After HAART's large innovation in efficacy in the mid-1990's, new drugs were less likely to be improvements on the efficacy dimension. Indeed, after the mid-1990's, average CD4 count rises to healthy levels, but stays below 600 (relative to 1000, which is roughly the average for HIV− people). This means that once products were developed that allowed patients to recover healthy (but not uninfected-level) CD4 counts, there is instead a rightward shift as innovations lead to reductions in side effects without noticeable improvements in efficacy. Consumer demand patterns suggest a preference for drugs with fewer side effects — especially when survival is less of a concern. The path of innovation seems to have followed this pattern. This is not surprising given that profit maximizing firms would presumably direct their inventive activities towards products where potential demand is high. In the following section, we formalize this idea, specifying a model where innovation is endogenous to aggregate consumer demand.

3 Model

This section describes a model of innovation in the market for HIV treatments. Each treatment is multidimensional: it can improve health, but can also have side effects that lead to physical ailments. Health and side effects can have both short-run and long-run consequences, affecting utility, labor market outcomes, future health and survival. New products are developed in clinical trials and both the entry of new products along with the exit of incumbent products are governed by stochastic processes that are endogenous to aggregate

[16] An exception is AZT, which remained a standard component of HAART.

consumer demand. The entry and exit processes we specify capture, in reduced form, how firms and government institutions (e.g., the FDA) interact to develop and introduce new medical treatments.

The data set we use provides rich information on consumers and so we model demand in greater detail than supply. Agents in the model maximize lifetime utility by choosing an HIV medical treatment. They can choose a product that is available on the market, access an experimental treatment by participating in a clinical trial, or opt for no treatment at all. Conditional on consumer characteristics, all treatments on the market cost the same. When making treatment decisions, consumers face several sources of uncertainty. They are uncertain about current-period outcomes, including their income and realized treatment side effects. They are also uncertain about the evolution of other individual-specific state variables, including their health, which affects future ailments and survival. Finally, consumers face uncertainty over the evolution of the product market resulting from entry and exit of products.

We highlight two features of the model. First, innovation is explicitly tied to consumer behavior. New products are drawn from a distribution that is a function of the share of patients in clinical trials along with endogenous market shares, both of which aggregate dynamically optimal individual choices. Second, when making treatment choices, consumers form beliefs over evolving choice sets arising from innovation and the exit of older products. In modeling consumer beliefs over future innovations, we avoid two simplifying assumptions. We do not assume that consumers have perfect foresight over future products. Nor do we assume that they are fully unaware of the innovation process, in which case changes to the choice set would amount to regime shifts. Rather, we model consumers as forming beliefs using the same stochastic processes governing entry and exit that we specify to capture supply. To understand these model features, consider the following example. A relatively healthy consumer may avoid choosing effective drugs with strong side effects in the current period if he believes that the introduction of effective drugs with fewer side effects is imminent. In contrast, a sick consumer may not want to avoid medication despite side effects if he fears that he may not survive until better drugs are introduced. Moreover, if sicker patients' choices accelerate innovation, this can further incentive healthier patients to delay using treatments.

Section 3.1 discusses the supply of treatment, including entry of new products and exit of incumbent products from the market. Section 3.2 specifies consumer demand for treatment, including choice sets, utility and individual state-to-state transitions.

3.1 Supply

We specify a reduced-form model of supply that captures the evolution of product characteristics. We do not model firm behavior, strategic pricing, R&D decisions or the role of government institutions.[17] Entry and exit occur at the end of the period immediately before the next period begins. We start by describing the aggregate state at the beginning of period t, followed by specifications for product entry and exit.

3.1.1 The Aggregate State

The aggregate state is denoted Ξ_t and summarizes market-level quantities at period t. It contains current and lagged product qualities $\mathbb{P}_t$, market shares $\mathbb{S}_t$, and the distribution of consumer characteristics $\mathcal{F}_t$, each described below. The aggregate state is given by:

$$\Xi_t = \{\mathbb{P}_t, \mathbb{S}_t, \mathcal{F}_t\}$$

Product Characteristics: Let $\mathcal{P}_t$ be the characteristics of products available at t. $\mathbb{P}_t$ denotes the set of characteristics for treatments available on the market in periods t, $t-1$ and $t-2$:

$$\mathbb{P}_t = \{\mathcal{P}_t, \mathcal{P}_{t-1}, \mathcal{P}_{t-2}\}$$

Market Shares: Let s_{t-1} denote a set containing the shares of products available at $t-1$. $\mathbb{S}_t$ is a set of product market shares for periods $t-1$ and $t-2$.

$$\mathbb{S}_t = \{s_{t-1}, s_{t-2}\}$$

Consumer Characteristics: $\mathcal{F}_t$ is the current distribution of consumer characteristics.[18]

3.1.2 Entry

We now explain product entry for new products that are available starting at the beginning of period $t+1$. The entry process is conditional on information that is in the state space at the start of period t. Entry of new products is modeled as a two-dimensional conditional distribution of new product characteristics $F_{\theta|\omega_t}$ conditioned on an endogenous reference point for innovation or *centroid*, denoted ω_t along with a distribution of number of new products

[17] Modeling supply in this way limits the sorts of counterfactuals we can perform. For example, our model would be ill-equipped to evaluate policies affecting market structure. We return to this point when discussing the counterfactual policy simulations we perform with the estimated structural model.

[18] The initial distribution of consumer characteristics is denoted $\mathcal{F}_0$.

F_N. Each of these objects is described in greater detail below.

Centroid: At any period t, the centroid for innovation ω_t is a two-dimensional weighted average of product characteristics for treatments available on the market in period $t-1$, where the weights are endogenous market shares:

$$\begin{aligned} \omega_t &= f_1^S(\mathbb{S}_t, \mathbb{P}_t) \\ &= \sum_{k \in \mathcal{P}_{t-1}} s_{kt-1}\theta_k. \end{aligned} \tag{1}$$

Each HIV treatment has two characteristics: its effectiveness at raising CD4 count, which we denote θ_k^h, and its propensity to cause side effects, denoted θ_k^x. The characteristics of product k are collected into a vector denoted $\theta_k \in \mathbb{R}^2$. Market share for product k is denoted s_{kt-1}, defined as the ratio of individuals who consume treatment k relative to the number of individuals who consume any treatment. The centroid essentially summarizes the evolving state of the product characteristics, serving as a baseline around which new products emerge.[19]

Characteristics of Trial Products and New Products: New products are developed in trials. Hence, the characteristics of the trial product at t as well as the characteristics of every new product introduced at $t+1$ are derived from the same process. The characteristics θ_{kt+1} of product k available at $t+1$ (either trial or newly introduced) are defined as an innovation around the previous period centroid as follows:[20]

$$\theta_{kt+1} = \left(\omega_{t+1}\mathbf{1}\{k = trial\}_{t+1} + \omega_t\left(1 - \mathbf{1}\{k = trial\}_{t+1}\right)\right) + \nu^*_{kt+1} \tag{2}$$

where $\mathbf{1}\{k = trial\}_{t+1}$ is an indicator for whether product k is trial or newly introduced at $t+1$. Notice that the difference between the trial product at $t+1$ and the newly introduced products is that the latter are innovations around last period's centroid. The magnitude and the direction of the innovation depend on the previous level of experimentation in clinical trials:

$$\nu^*_{kt+1} = \Theta_0^\nu + \Theta_1^\nu \cdot TrialsShare_t + \nu_k \tag{3}$$

[19] If nobody uses a treatment the base for innovation remains the same, i.e. $\omega_t = \omega_{t-1}$.

[20] In estimation, we tested alternative specifications of equation (2) in which the characteristics of the previous trial product determine the characteristics of new products. Although these specifications are intuitive if, for instance, better trial products lead to better new market products, the relation between current trial product characteristics and the characteristics of future market products is statistically insignificant.

In equation (3), Θ_0^ν and Θ_1^ν are two-dimensional vectors of parameters. Θ_1^ν captures how participation in clinical trials can shift the magnitude and the direction of innovation.[21] The share of individuals who consumed a trial product in t is denoted $TrialsShare_t$, which is an endogenous object derived from the consumers' decision problems. To express this point, we write it as a function of product and consumer characteristics as follows:

$$TrialsShare_t = f_2^S(\mathbb{P}_t, \mathcal{F}_t). \tag{4}$$

The size of the innovation also depends on ν_k, a two-dimensional vector of mean-zero exogenous disturbances in technology drawn from a two-dimensional stationary distribution, denoted F_ν. We do not specify a parametric form for F_ν. As will be explained in Section 4, when we discuss estimation, F_ν is a non-parametric distribution estimated using the full history of observed innovations around the centroid.

Using equations (1)-(3), we are now able to define the distribution of new product characteristics θ_{k1+1}, denoted $F_{\theta|\omega_t}$. The distribution is a translation of F_ν centered on $(\omega_t + \Theta_0^\nu + \Theta_1^\nu \cdot TrialsShare_t)$.

Number of New Products: We now specify the process governing the number of new products drawn from $F_{\theta|\omega_t}$ and which enter the market at $t+1$. The number of new products is denoted New_{t+1} and is distributed according to F_N. We specify F_N as a negative binomial that permits dispersion in the mean:

$$\begin{aligned} New_{t+1} &\sim Poisson\,(\mu_t^*) \\ \mu_t^* &\sim Gamma\left(1/\alpha^N, \alpha^N \mu_t\right) \\ \mu_t &= \exp(\beta_0^N + \beta_1^N Q_t + \beta_2^N TrialsShare_{t-1}) \\ \ln \alpha_t^N &= \alpha_0^N + \alpha_1^N Q_t \end{aligned} \tag{5}$$

The binomial model is conditioned on the magnitude of innovations introduced at time t, denoted Q_t, along with trial participation during period $t-1$. The inclusion of the trial share is motivated by empirical evidence showing that more experimentation can be conducted if larger proportions of consumers participate in clinical trials. We include Q_t to capture the fact that large breakthroughs tend to be followed by a relatively large number of new products. This may occur if breakthroughs spur innovative activity as firms attempt to capture market share. The magnitude of previous innovations measures the distance (in

[21] For instance, larger parameter values imply that increased trial participation leads to larger innovations on average. Alternatively, parameter values could also imply that higher levels of trial participation lead to larger effectiveness versus side effects innovations if the parameter mapping $TrialShare_t$ to effectiveness is larger than the parameter mapping it to side effects.

characteristics space) between products at time t and the previous period centroid around which they were drawn:[22]

$$\begin{aligned} Q_t &= f_3^S(\mathbb{S}_t, \mathbb{P}_t) \\ &= \sum_{r \in \{h,x\}} \frac{\max_{\theta^r \text{ new at } t} \left\{\theta^r - \omega_{t-1}^r\right\}}{\max_{\theta^r \text{new at } \tau, \forall \tau} \left\{\theta^r - \omega_{\tau-1}^r\right\}} \end{aligned} \tag{6}$$

In summary, ω_t, $F_{\theta|\omega_t}$, and F_N imply that the path of innovation is endogenous to consumer demand. Individual choices, aggregated into market shares, affect the centroid in equation (1). By affecting ω_t, market shares affect the characteristics of each new product θ_{kt}, as well as the future trial product, according to equation (2). Intuitively, treatments that keep patients alive and those associated with fewer ailments capture larger shares of the market. These more popular treatments hold greater weight in the centroid around which innovations are drawn. Additionally, aggregate participation in clinical trials affects both the speed and the direction of innovation, captured by equation (3). Finally, both trial participation and the magnitude of previous innovations affect the expected number of new products introduced into the market.

3.1.3 Exit

Exit of incumbent treatments happens at two different levels: *exit for switchers* and *overall exit*. *Exit for switchers* happens when the product is no longer available for individuals who have yet to use it, but is still available for those who were consuming the product in the prior period. *Overall exit* happens when the product is no longer available to any consumer. Exit happens according to the following rules that aim to reconcile empirical observations and theory.[23]

1. If the ratio of people switching to product k relative to the total number of individuals switching onto a new product falls bellow $\tilde{\sigma}_1$ during three consecutive periods, the product is withdrawn from the market. $\tilde{\sigma}_1$ is chosen as the minimum conditional share observed in the data and the number of consecutive periods (three) is chosen so that a single period of low demand does not lead to a premature exit.

[22] The relative change is computed for each of the two dimensions of product characteristics (health and lack of ailments) and is scaled by the maximum change observed over the sample period. In order to compute Q_t we need the scaling quantities $\max_{\theta^r \text{ new at } \tau, \forall \tau} \left\{\theta^r - \omega_{\tau-1}^r\right\}$ for $r \in \{h, x\}$ which are estimated consistently by their data counterparts.

[23] Expected shares must be positive due to model assumptions on the taste shocks of consumers. Distributional assumptions on taste shocks are stated below, when we discuss the demand portion of the model.

2. If the ratio of people consuming product k (either by staying or by switching onto it) relative to the total number of people consuming any market product falls below $\tilde{\sigma}_2$ during two consecutive periods, the product is withdrawn from the market. $\tilde{\sigma}_2$ is chosen as the minimum conditional share observed in the data.

The exit criteria can be written in terms of the aggregate state of the problem as follows:

$$ProductsWithdrawn_{t+1} = f_4^S(\mathbb{S}_t, \mathbb{P}_t, \mathcal{F}_t)$$

3.1.4 The Evolution of the Aggregate State

Given the current aggregate state Ξ_t and the exogenous distribution of innovations F_ν, aggregate choices induce a new distribution of consumer characteristics $\mathcal{F}_{t+1}$. Through the entry and exit mechanisms, a new set of available products comes available and is denoted $\mathcal{P}_{t+1}$, which is used to form $\mathbb{P}_{t+1}$. Finally, consumer choices can be summarized into market shares $\mathcal{S}_t$, which are used to form $\mathbb{S}_{t+1}$. Thus, we have all the components of the one-period-ahead aggregate state Ξ_{t+1}, which captures the supply of medical treatment. We now turn to consumer demand.

3.2 Demand

The individual chooses medical treatment to maximize expected discounted lifetime utility. In making decisions, he observes his current state which includes individual-specific variables, such as health, along with market-level variables, such as the current state of medical technology. Individuals use market-level variables to form expectations over the future path of innovation. In specifying the individual's problem, we discuss state variables, the choice set, flow utility and stochastic processes governing outcomes and state-to-state transition probabilities. Next, we discuss consumer information and aggregate state forecasts. We conclude this section by specifying the value function.

3.2.1 State Variables

The state for individual i at period t is denoted $\mathcal{Z}_{it}$, where

$$\mathcal{Z}_{it} \equiv \langle z_{it}, \varepsilon_{it} \rangle \tag{7}$$

z_{it} is a set of state variables that is further sub-divided into a set of individual-specific variables, denoted $z_{it}^{\mathcal{I}}$, and a set of aggregate variables denoted $z_t^{\mathcal{M}}$:

$$z_{it} \equiv \left\langle z_{it}^{\mathcal{I}}, z_t^{\mathcal{M}} \right\rangle \tag{8}$$

The individual-specific state variables, $z_{it}^{\mathcal{I}}$, are

$$\begin{aligned}
b_i &: \text{a set of race indicators} \\
edu_i &: \text{a set of time invariant education indicators} \\
h_{it-1} \in \mathbb{R}_+ &: \text{health at the start of period } t \\
a_{it-1} \in \{25, 25.5, \ldots\} &: \text{age at the start of period } t \\
l_{it-1} \in \{0, 1\} &: \text{worked during the prior period } t-1 \\
q_{it-1} = \left\{q_{it-1}^x, q_{it-1}^h\right\} \in \mathbb{R}^2 &: \text{characteristics of product consumed last period} \\
\eta_i &: \text{person-specific income characteristic}
\end{aligned}$$

The individual is either white, black or Hispanic and has one of four mutually exclusive educational categories: high school, some college, college or more than college. His health, measured by CD4 count, is a continuous positive number.[24] His age is measured in half-year increments, corresponding to the frequency of MACS data collection. l_{it-1} indicates whether the individual worked last period. If the individual consumed a market product in the prior period, the characteristics of that product, denoted q_{it-1}, are part of his current state space. q_{it-1}^x measures lack of side effects of the treatment and q_{it-1}^h measures treatment effectiveness. η_i is an exogenous person-specific characteristic that affects his income generating potential. Besides individual-specific variables, z_{it} contains aggregate level components, collected in $z_t^{\mathcal{M}}$, which individuals use to forecast the evolution of the market. $z_t^{\mathcal{M}}$ will be described further in Section 3.2.5. Individuals also receive a vector of choice-specific additive utility disturbances ε_{it} which are assumed independent across time, individuals and choices.[25]

3.2.2 Choices

At each period t the individual chooses whether or not to use medication. If he opts for medication, he may choose the same product he consumed in the previous period or he may choose from the set of other treatments currently available on the market. Alternatively, he may choose a trial treatment. The individual faces uncertainty about the quality of both market and trial treatments.

[24] CD4 ranges from 0 to 2915 in our analytic sample with a median of 448. Healthy CD4 counts are those above 500 units per mm^3 and typically range between 500 and 1,500.

[25] In estimation, η_i and ε_{it} will be unobserved to the econometrician.

Market Products: We begin by describing how consumers choose market products. The individual learns about the quality of a product immediately after using it. Hence, if he chooses the same market treatment he consumed last period, he faces no uncertainty regarding its characteristics.[26] Alternatively, if he decides to try a different market drug, his alternative is to choose one among several groups or *clusters*, which contain drugs with similar qualities. The agent is then randomly assigned a drug within the cluster he selected.

Formally, at every period t there is a set of market products $\mathcal{P}_t$ clustered in several groups collected in $\mathcal{G}_t$ with typical element g_t. $\mathcal{G}_t$ denotes both the collection of clusters available at t and the cardinality of the collection (i.e., the total number of clusters). When individual i decides to consume a market treatment that is different from the one he consumed last period, he must choose a cluster $g_t \in \mathcal{G}_t$. By selecting g_t, he chooses random assignment to one of all products in g_t. Our clustering process is a device to make the model tractable and estimation feasible by substantially reducing the state space while still allowing individuals to choose among numerous medical treatments. This approach also captures the idea that consumers often choose drugs after observing product labels without knowing specific drug characteristics beyond the fact that certain labels are associated to a particular mean and variance of characteristics. Product clusters at t yield from a k-means algorithm (see Appendix B). At any given period we set the maximum value of $\mathcal{G}_t$ at $\mathcal{G}^{\max} = 3$ so that the individual knows how many groups will be available every period. Our choice of $\mathcal{G}^{\max}$ guaranties that there is a non-negligible number of consumers choosing each cluster.

Conditional on choosing a cluster, the probability of being assigned one of the products within the cluster is given by weights that depend on the treatment characteristics and the number of products in the cluster. The weight of product k in cluster g_t is given by the following equations

$$\tilde{s}_{k|g_t} = \frac{\mathbb{E}\left[s_{k|g_t}|X^w_{k,t}\right]}{\sum_{r\in g_t}\mathbb{E}\left[s_{r|g_t}|X^w_{k,t}\right]} \tag{9}$$

$$\mathbb{E}\left[s_{k|g_t}|X^w_{k,t}\right] = \exp\left(X^w_{k,t}\beta^w\right) \tag{10}$$

where $X^w_{k,t}$ includes a constant term, the ranking (within its cluster) of the characteristics of the product, the number of members in the cluster, whether the product is new, and several interactions. We assume that agents do not observe the characteristics of each individual product. Rather, consumers at t observe the first two moments of the cluster distribution implied by the weights. These cluster-level characteristics observed by the consumer are

[26] As discussed above, his state space includes the characteristics of the drug consumed in the prior period q_{it-1}.

denoted

$$W_t = f_1^D(\mathcal{P}_t). \tag{11}$$

Trial Products: If the individual chooses neither a cluster nor to stay in his previous treatment, he may instead join a clinical trial to get an experimental treatment. Trial product characteristics are unknown, but are draws around the centroid ω_{t-1} according to equation (2).[27] Thus, consumers are aware that product characteristics of trial treatments are distributed according to $F_{\theta|\omega_t}$, which is equivalent to the distribution of new product characteristics. A key difference between choosing a cluster of products g_t and the trial treatment is that, after choosing group g_t and once the quality of the assigned product is learned, the consumer has the chance of choosing the same treatment with certainty the next period. In the case of a trial, the consumer must always take a new draw from the innovation distribution.

Having described each option, we now formally specify the choice set. Let d_{jit} be the choice indicator that takes the value of one if agent i in period t chooses medical treatment j in the choice set $\mathcal{C}_{it}$. The choice set is time-specific because the characteristics of available products evolve with entry and exit of products. The choice set is also individual specific since individuals who chose a market treatment in the prior period may choose that treatment again. If the individual did not choose a market treatment in the prior period his choice set is:

$$\mathcal{C}_{it} = \begin{cases} 0 & \text{No Treatment} \\ 1 & \text{Cluster } g_t = 1 \\ 2 & \text{Cluster } g_t = 2 \\ \vdots & \vdots \\ \mathcal{G}^{max} & \text{Cluster } g_t = \mathcal{G}^{max} \\ \mathcal{G}^{max} + 1 & \text{Trial} \end{cases} \tag{12}$$

If the individual chose a market treatment in the prior period, his choice set $\mathcal{C}_{it}$ is augmented by one alternative to include the possibility of consuming his previous period treatment with full knowledge of its characteristics.[28]

[27] One way to think about this point is that consumers entering a trial see ω_{t-1} as the quality of a placebo drug administered in a trial. This makes sense in the context of HIV since new drugs are not tested against no drug at all, but are instead tested against "current best practices" (see Ickovics and Meisler (1997)).

[28] To clarify, if a consumer chooses a cluster in period t, in period $t+1$ he may choose the treatment he was randomly assigned to in period t. Alternatively, he may choose the same cluster, which means he is randomly assigned once again to a treatment in the cluster.

3.2.3 Utility

For choice $j \in \mathcal{C}_{it}$ and state z_{it}, the utility at period t for individual i is a function of his health, ailments and net income, given by

$$y_{jit} + \varepsilon_{jit} = \alpha_m(m_{jit} - o_{jit}) + \alpha_{jit}\left(z_{it}\right) + \alpha_{xp}x_{jit}d_{0it} + \varepsilon_{jit} \tag{13}$$

The first expression on the right-hand-side of equation (13) is gross income m_{jit} minus out-of-pocket payments for medical treatment o_{jit}, so that α_m captures consumption utility. The second expression captures choice-specific utility. $\alpha_{jit}\left(z_{it}\right)$ is the sum of choice-specific preference parameters that depend on race, age and health and is defined as follows:

$$\alpha_{jit}\left(z_{it}\right) \equiv \alpha'_{jb}b_i + \alpha_{ja}a_{it-1} + \alpha_{jh}h_{it-1} \tag{14}$$

For clusters, market treatment demand is modeled using a characteristics approach, which means that we omit cluster-specific dummy variables in the utility function and instead assume that variation in cluster choices is fully captured by cluster characteristics (effectiveness and side effects). This implies that parameters α'_{jb}, α_{ja}, and α_{jh} are constant across clusters. The utility function also includes separate sets of parameters for clinical trials participation and for staying on a current treatment in the next period. This captures how experimentation in treatment choices can imply additional costs or benefits. In the case of clinical trials, utility parameters capture, for example, the fear of trying an experimental drug or preferences for altruism since trial participation may help future patients. In the case of continuing to use the same product, utility parameters may capture a preference for certainty, which could help to explain consumer reluctance to switch even when better products enter the market.

Health affects lifetime utility through its impact on future health and survival. It also affects the probability of suffering physical ailments, which affect utility directly and through earnings. These processes are described in Section 3.2.4. In addition, we allow health to directly affect flow utility by interacting it with the indicator variable for trial participation. This captures the possibility that the time and psychic costs of finding a trial slot can vary by health if, for example, doctors are more willing to encourage experimentation or if trial slots are more readily available for sicker patients. Similarly, we interact lagged health with choosing a cluster to capture how individuals may be more willing (or encouraged by doctors) to choose random assignment to a cluster treatment when in poor health. We normalize α_{jh} to zero for those who continue using the same product. Therefore, α_{jh} captures the additional effect that health has on the utility of experimenting with a new treatment, either through

cluster use or trial participation.

The third expression in equation (13) captures the utility cost of physical ailments. $x_{jit} = 1$ indicates that the individual does not suffer from physical ailments and d_{0it} is an indicator for not consuming a medical treatment. The interaction captures how distaste for ailments can vary depending on whether or not a treatment is being consumed. For example, distaste for physical ailments may be less strong for individuals using a treatment, in which case ailments are a sign of effectiveness. We normalize the utility cost of ailments while using a treatment to zero. Hence, α_{xp} represents the differential distaste for ailments for individuals who are not taking a treatment.

The fourth expression in (13) is ε_{jit}, which are unobserved choice-specific taste shocks. They are Extreme Value Type I distributed and are assumed independent across choices, individuals and time. Finally, for identification purposes, we normalize the non pecuniary benefits from not consuming a treatment to zero (Magnac and Thesmar, 2002; Arcidiacono and Miller, 2015)

3.2.4 Outcomes and Transitions

In this section, we specify the stochastic processes governing the evolution of state variables in z_{it} as well as the outcome variables: income, out-of-pocket payments, ailments, and survival.

Income: Gross income is a function of today's state, z_{it}, and ailments, x_{jit}. It is given by

$$m_{jit} = X^m_{jit}\Gamma^m + \eta_i + \epsilon^m_{it} \tag{15}$$

where $X^m_{jit} = [1, h_{it-1}, \ldots, h^7_{it-1}, a_{it-1}, a^2_{it-1}, b_i, edu_i, l_{it}, x_{jit}]$ and η_i is the individual-specific income generating potential. Gross income does not include product cost, which is accounted for in the payments equation below. Individuals observe the income iid shocks ϵ^m_{it} before making their treatment choice.

Payments: Out-of-pocket payments are censored at zero. They are given by the following tobit specification

$$o_{jit} = o\left(X^o_{jit}, \epsilon^o_{it}; \Gamma^o\right) \tag{16}$$

where $X^o_{jit} = [1, h_{it-1}, \ldots, h^6_{it-1}, a_{it-1}, a^2_{it-1}, b_i, edu_i, \{d_{jit}\}^5_{j=0}, l_{it}, x_{jit}]$ and ϵ^o_{it} is the error term in the underlying equation. Since we do not directly observe prices, and in order to simplify the problem, we assume a constant cost of participating in a trial as well as a constant cost of consuming a market product.[29]

[29]End-users customarily pay a standardized deductible that is a fraction of the brochure price of the drug paid by the insurance company. Median out-of-pocket drug costs are about $300 every six months for a

Labor Supply: We do not model labor supply explicitly as a choice as it is not the main purpose of this paper. However, labor supply may be affected by treatment choices, e.g., through health status and physical aliments. Moreover, labor supply affects income and therefore utility. To capture this, we treat labor supply as a state variable that individuals know at the beginning of the period before making their treatment decision. Individuals draw their labor market participation from the distribution characterized by

$$\Pr[l_{it} = 1|X_{it}^l] = \frac{1}{1 + \exp(X_{it}^l \Gamma^l)} \tag{17}$$

where $X_{it}^l = [1, l_{it-1}, h_{it-1}, \ldots, h_{it-1}^4, a_{it-1}, a_{it-1}^2, b_i, edu_i]$.

Physical Ailments: First, define a mapping from the choice to the characteristics of the treatment

$$\theta(d_{jit}) = \{\theta^x(d_{jit}), \theta^h(d_{jit})\} \tag{18}$$

where $\theta(d_{jit}) = q_{it-1}$ if the individual consumes his prior-period market treatment. $\theta(d_{jit})$ is a stochastic variable if the individual chooses a cluster or if he joins a trial. A production function transforms drug characteristics and health into aliments. Let x_{jit} be an indicator that takes the value of 1 if the individual does not suffer ailments in t after choosing alternative $j \in \mathcal{C}_{it}$. The probability of not having physical ailments for individual i is modeled as:

$$\Pr\left[x_{jit} = 1|\cdot\right] = \frac{\exp\left(\sum_{m=0}^{5} \gamma_m^x h_{it-1}^m + \theta^x\left(d_{jit}\right)\right)}{1 + \exp\left(\cdot\right)} \tag{19}$$

Health: CD4 count is our objective measure of health. Like ailments, health at the beginning of period t is a function of previous health and drug characteristics. The production function of health is specified as:

$$h_{jit} = \sum_{m=0}^{5} \gamma_m^h h_{it-1}^m + \theta^h\left(d_{jit}\right) + \epsilon_{it}^h \tag{20}$$

We assume that $\mathbb{E}[\epsilon_{it}^h|\cdot] = 0$, where the expectation is conditional on the vector of regressors of the health production function.[30]

Survival: At the end of any period t individuals may survive into the next, denoted by $S_{it+1} = 1$, with the following probability

$$D_{it+1}\left(z_{it+1}\right) \equiv \Pr[S_{it+1} = 1|z_{it+1}] = \frac{1}{1 + \exp(X_{it}^d \Gamma^d)} \tag{21}$$

regime of drugs that would cost the insurance company between \$5,000 and \$15,000.

[30] We do not make parametric assumptions on the health disturbance. It is estimated non-parametrically using the residuals of the health production function.

where $X_{it}^d = [1, h_{jit}, \ldots, h_{jit}^5, a_{it}, a_{it}^2, b_i, edu_i, x_{jit}]$.

3.2.5 Consumer Information and Aggregate State Forecasts

We assume that consumers have rational expectations, but they do not observe the entire aggregate state Ξ_t. Instead, they observe the reduced aggregate state $z_t^{\mathcal{M}}$, which is a mapping from Ξ_t, and integrate over the treatment characteristics that they do not observe. The aggregate portion of the individual's state is given by

$$z_t^{\mathcal{M}} \equiv \langle \omega_t, W_t, \mathcal{F}_t \rangle \tag{22}$$

The individual observes the centroid for innovation ω_t, described in Section 3.1, which determines the expected characteristics of trial products. His information set also contains the characteristics W_t of the clusters of products he observes, described in Section 3.2. Finally, he observes the current distribution of consumer characteristics $\mathcal{F}_t$. When agents form expectations over future innovations, they are likewise assumed to observe cluster characteristics rather than the characteristics of each treatment. A timeline of the model can be found in Appendix B.

3.2.6 The Value Function

We define the value function conditional on choice $j \in \mathcal{C}_{it}$, net of taste shocks, for individual i at time t as follows:

$$v_{jit}(z_{it}) = \mathbb{E}_y[y_{jit}|z_{it}] + \beta \mathbb{E}_z \left[D_{it+1}\left(z_{it+1}\right) \mathbb{E}_\epsilon \left[\max_{c \in \mathcal{C}_{it+1}} \left\{ v_{cit+1}(z_{it+1}) + \varepsilon_{cit+1} \right\} \right] \middle| z_{it}, j \right] \tag{23}$$

Expectations are taken over product characteristics affecting the flow utility and the evolution of both observed and unobserved state variables. The first expectations operator, $\mathbb{E}_y$, denotes expectations over outcomes that affect flow utility, including income and physical ailments. The second operator, $\mathbb{E}_z$, denotes expectations over the evolution of observed state variables z_{it}. The third operator, $\mathbb{E}_\epsilon$, denotes expectations taken over the joint distribution of future unobserved choice-specific taste shifters. The value function is fairly standard except for the time and individual subscripts on the choice set $\mathcal{C}_{it+1}$. These subscripts capture the evolving choice set due to product innovation, which makes the problem non-stationary and therefore not estimable using standard methods typically used to evaluate conditional value functions. We now turn to a discussion of how we estimate the structural model.

4 Estimation

We use GMM to estimate model parameters. In Section 4.1, we provide an overview of the estimation procedure, summarizing the algorithm, which includes 5 steps. The first four steps constitute the "first stage," used to obtain quantities that do not change with utility parameters, including simulation of future choice and technology paths. We compute these quantities a single time and use them in the "second stage" to construct moments used in GMM estimation of utility parameters. In Section 4.2, we provide further details on the GMM estimator. We first describe the theoretical moment conditions and their sample analogs. We also discuss our forward simulation procedure, which is designed to incorporate non-stationarity, endogenous innovation. Finally, we discuss estimation of the non-parametric distribution of innovations. A more extensive treatment of the estimation procedure is found in Appendix C.

4.1 Overview

Our estimation procedure can be summarized in the following steps:

1. *Products.* We start by defining a product as a combination of single-product components. AZT or the combination of AZT+3TC+Saquinavir are both examples of products in our framework. We define one single trial product per period as the one used by those individuals joining a clinical trial. Given this definitions, we estimate product characteristics together with the health and no-ailment processes (in equations (19) and (20)) using equations (S5) and (S6) in Appendix C.[31]

2. *Clusters.* Using the estimated product characteristics in step 1, we use a k-means algorithm to obtain clusters of products for every period (see equation (S1) in Appendix B). Then, using the characteristics of the products in each cluster, we obtain within cluster weights for each product in each cluster (see equations (9) and (10)). Finally, using within cluster weights we compute cluster characteristics—mean and variance matrix.

3. *Innovation.* We back out centroids for innovation for each period (see equation (1)) using product characteristics from step 1. Then, since every product (new and trial) is modeled as a draw around the centroid (see equation (2)), for every new and trial product at a given period we compute the realized innovation as the residual from

[31] Product characteristics are obtained using product indicators in the estimating equations (S6) and (S5) in Appendix C.

subtracting the relevant centroid from the product characteristic (step 1). Using the realized innovations we estimate the parameters of the innovations equation (3). Next, we non-parametrically estimate the stationary distribution of innovations F_v using the residuals ν_{kt+1} from estimating equation (3). Finally, we use the number of new products per period to estimate the distribution of number of new products (see equations (5) and (6)).

4. *Outcomes.* We estimate processes for income, out-of-pocket payment, labor supply and survival (see equations (15), (16), (17), and (21)).

5. *Utility Function.* We estimate the utility parameters in equation (13) using a GMM estimator and moment conditions that equate two alternative representations of differences in conditional value functions, one based on current conditional choice probabilities (CCPs) and the other based on future CCPs and simulated continuation values. In order to obtain these moments we estimate flexible parametric CCPs using cluster characteristics from step 2, centroids from step 3 and other aggregate and individual-specific state variables (see Appendix C). We then use forward simulation to generate choice and technology paths as well as future individual states that serve as inputs to the simulated continuation value. In our forward simulation we use the estimated CCPs as well as all estimated results from steps 1 through 4.

4.2 Moment Condition

Our moment conditions rely on differences between alternative representations of the difference in conditional value functions $v_{jit}(z_{it}) - v_{oit}(z_{it})$.[32] The first represesntation is the log odds ratio formed with current-period conditional choice probabilities. The second representation relies on the results in Proposition 1, which yields the conditional value function as a mapping of future conditional choice probabilities and utility parameters.

Proposition 1. *Let $V(z_{it}, \varepsilon_{it})$ be the value function for individual i at period t who has a state given by z_{it} and ε_{it}. Let $p_{jit}(z_{it})$ be the probability that individual i chooses option j at time t conditional on his state z_{it}. Define $\mathbb{E}[\,|\, j, d_i^o]$ as the expectation conditional on decision j at t and optimal behavior, denoted d_i^o, up to some period $T^* > t$.*[33] *Define $\psi_{kit}(z_{it}) \equiv \mathbb{E}_\varepsilon[\varepsilon_{kit} | d_{it}^o = k, z_{it}]$ as the expected value of the kth taste shock conditional on alternative k being optimal. Finally, let γ be the Euler constant. Then, the conditional value*

[32] These moment conditions appeal to well-known results following from our assumption that the taste shocks ε_{jit} are iid Extreme Value Type I distributed (Hotz and Miller (1993)).

[33] Recall that D_{it+s} is the one-period-ahead probability of survival, defined in Section 3.2

function can be written as

$$
\begin{aligned}
&v_{jit}(z_{it}) = \mathbb{E}_y[y_{jit}|z_{it}] \\
&+\sum_{s=1}^{T^*} \beta^s \mathbb{E}\left[\left(\prod_{r=1}^{s} D_{it+r}\left(z_{it+r}\right)\right) \sum_{k \in C_{t+1}} p_{kit+s}\left(z_{it+s}\right)\left[y_{kit+s}\left(z_{it+s}\right) + \psi_{kit+s}\left(z_{it+s}\right)\right] \middle| z_{it}, j, d_i^o\right] \\
&+\beta^{T^*+1} \mathbb{E}\left[\left(\prod_{r=1}^{T^*} D_{it+r}\left(z_{it+r}\right)\right) V\left(z_{it+T^*+1}, \varepsilon_{it+T^*+1}\right) \middle| z_{it}, j, d_i^o\right]
\end{aligned} \tag{24}
$$

and

$$
\psi_{kit}\left(z_{it}\right) = \gamma - \ln\left(p_{kit}\left(z_{it}\right)\right) \tag{25}
$$

Proof: see Appendix C

Let $J = 6$ be the maximum possible cardinality of the individual's choice set and let $w\left(z_{it}\right)$ be a vector of instruments orthogonal to the difference between alternative representations. Using equation (24) we can form the following moment conditions:

$$
\mathbb{E}\left\{ w\left(z_{it}\right) \otimes \begin{bmatrix} \ln\left(\frac{p_{oit}(z_{it})}{p_{1it}(z_{it})}\right) + v_{1it}(z_{it}) - v_{oit}(z_{it}) \\ \vdots \\ \ln\left(\frac{p_{oit}(z_{it})}{p_{J-1it}(z_{it})}\right) + v_{J-1it}(z_{it}) - v_{oit}(z_{it}) \end{bmatrix} \right\} = 0. \tag{26}
$$

4.2.1 Sample Analog and Forward Simulation

To form sample analogs of the moments in equation (26), we first substitute the theoretical log odds ratio using the estimated CCPs. Second, we use Proposition 1 to obtain differences in conditional value functions using forward simulation (Hotz et al., 1994). In our forward simulation procedure, for every individual i at time t facing choice set $\mathcal{C}_{it}$, we fix choice j and use the estimated stochastic processes governing outcomes and transitions to simulate his state variables at $t+1$. We then use the estimated parameters of the CCPs to simulate $t+1$ choices conditional on the new simulated state. We continue the same process until T^*, the value of which is set high enough so that the product $\beta^{T^*+1} \prod_{r=1}^{T^*} D_{it+r}\left(z_{it+r}\right)$ approaches zero, eliminating further differences in conditional value functions beyond T^*.

Forward simulation is used in a variety of settings to compute conditional valuation functions. The procedure we use is designed to capture how individuals making decisions form beliefs aware of the stochastic processes linking aggregate behavior to the evolution of choice sets. In particular, we first simulate aggregate behavior forward, which allows us to construct paths of technological innovation. We construct one such artificial path for each observation in the data. In other words, for every individual i at period t we forward

simulate the choices of all individuals in the sample at period t, and gather the technological path generated by their collective choices. Then, because individuals are atomistic, for each observation we can generate several sequences of future choices and payoffs taking as given a random subset of artificial technological paths.[34] This serves two purposes. It maintains the assumption, needed for consistency of the estimator, that the sample draws from the moment conditions—the contributions from each observation—are independent from each other. Additionally, using a random selection of artificial technological paths for every observation prevents simulation errors in technology paths from propagating across all observations.

The CCPs used in the forward simulation procedure condition on $z_t^{\mathcal{M}}$, which is the market-level information we assume consumers observe, rather than the full aggregate state Ξ_t. This captures how agents in the model make decisions using limited information about the aggregate state. In particular, we obtain individuals' expectations by simulating future paths of the full aggregate state Ξ_t using choices generated by individuals who make decisions with limited information $z_t^{\mathcal{M}}$.[35]

4.2.2 Estimating the Non-Parametric Distribution of Innovations F_ν

The forward simulation of future choice sets relies on the distribution of innovations. According to equations (2) and (3), the characteristics of trial products and of new products entering the market today are determined by the centroid (current or previous), previous trial participation and a draw from the distribution of innovations F_ν.[36] After computing centroids for innovation ω_t given by equation (1), for the trial product and for each new product at t, characterized by θ_{kt}, we compute a realized innovation vector as

$$\nu_{kt}^{\theta*} = \theta_{kt} - (\omega_t \mathbf{1}\{k = trial\}_t + \omega_{t-1}(1 - \mathbf{1}\{k = trial\}_t))$$

We do not impose that innovation vectors cannot be strictly negative. In other words, relative to the centroid, inferior products with lower quality in both dimensions (health

[34] We generate twenty sequences of choices and payoffs per observation.

[35] Because we simulate the future path of Ξ_t, we can also obtain the simulated future path of the subset $z_t^{\mathcal{M}}$.

[36] To estimate the innovation equation (3) and the distribution of innovations we use all periods in the MACS data with relevant information on treatment consumed, health and ailments (1986 to 2008). Over the time span in our data we observe 76 realized innovations from newly introduced market products and 22 realized innovations from trials products. Consistent with our definition of market products, we only consider trial products that entail at least 40 users.

and ailments) may be introduced.[37] We use the realized innovations vectors to estimate parameters Θ_0^ν and Θ_1^ν in equation (3). Using the residuals from this exercise, we estimate F_ν non-parametrically.

5 Parameter Estimates and Choice Dynamics

In this section we describe estimates that affect innovation (Section 5.1), present our estimated utility parameters (Section 5.2), and discuss parameters governing state-to-state transitions and outcomes (Section 5.3). Finally, in Section 5.4 we assess model fit.

5.1 Innovation

The innovation process is described in equations (2) and (3). Estimates of Θ_0^ν and Θ_1^ν from equation (3) are found in the top panel of Table 3. Recall, both parameters are two-dimensional vectors that map trial participation to the characteristics of new treatments. The first panel of Table 3 contains parameters relating trial participation to trial treatment and new treatment effectiveness and side effects. According to the estimates of Θ_1^ν, higher rates of consumer participation in clinical trials lead to improvements in new drugs on both dimensions of quality. However, the magnitude of this effect varies across quality dimensions. To see this, consider the differences in constants (Θ_0^ν), both of which are negative, suggesting that new drugs would on average be worse (relative to the centroid) if no consumers participated in trials. While expected health innovations are positive for lagged trial shares above 5.6%, the same is true for expected innovations on the ailments dimension when trial participation is 7.7%. As average trial participation in our sample is 7% (Table 2), if we ignored the role of the centroid, on average the quality of products would appear to improve over time in terms of health quality but would remain largely unchanged in terms of the ailments they cause through side effects. However, new products are drawn around the centroid, which is time-varying and a function of market shares. Therefore, if individuals prefer treatments with fewer side effects, the centroid will move in that direction of the characteristics space, which means more innovations will be drawn from that part of the space. In this sense, future products tend to be similar to products with larger market shares.

The stochastic component of an innovation is a draw from the non-parametric distribution F_ν showed in Figured 6. F_ν which is mechanically centered on $(0, 0)$, is unimodal.[38]

[37] This is not at odds with what we observe in the data, and theoretical reasons why this may happen have been provided in the literature (Miller, 1988).

[38] F_ν is centered on $(0, 0)$ because it is estimated from the residuals of equation (3). In a separate exercise,

Hence, conditional on trial participation, most innovations are small improvements. Because products are multidimensional, it is possible for the stochastic component to generate products that are more efficacious, but which cause worse ailments via side effects, or vice versa. However, according to the bottom panel of Table 3, which shows the covariance matrix of the distribution of innovations, there is a positive correlation of about 0.24 between the two quality dimensions of an innovation. Therefore, improvements to effectiveness tend to have fewer side effects than existing technologies once we have conditioned on trial participation rates.

Estimates of the distribution of the number of new products are shown in Table 4. Estimates suggest that large positive innovations in previous periods lead to larger numbers of new products. The magnitude of previous innovations also reduces the dispersion around the number of new products that enter. Both patterns are consistent with firms vying for market share following breakthroughs by producing similar products. The share of consumers opting to participate in a clinical trial in the prior period also increases the likelihood of more products entering the market. Our interpretation is that as more consumers select trial products firms increase their experimental activity. This leads to an increase in the quantity of viable new treatments that can be introduced into the market. The estimated distribution of the number of new products fits the data very well, according to Figure 7, which plots the empirical distribution along with the average (over time periods) of the predicted probabilities generated by the model.

5.2 Utility Parameter Estimates

Utility parameters are reported in Table 5. Individuals gain positive utility from income net of out-of-pocket treatment costs, which captures consumption utility. Moreover, a lack of physical ailments enters positively into the flow utility.[39] Prior literature has shown that even in the context of a deadly infection (HIV) individual treatment choices reflect a distaste for side effects (Chan and Hamilton, 2006; Papageorge, 2016). Recall, we have normalized to

we estimate the non-parametric distribution of innovations using directly the realized innovations from equation (2), ignoring equation in (3). This yields a bimodal distribution of innovations with one of the modes located approximately at the status quo point $(0, 0)$, and a second mode located north of the first one along the health axis. As shown in Figure 6, F_ν is no longer bimodal once we control for previous aggregate experimentation in clinical trials.

[39]Even though both net income and physical ailments parameters become insignificant once the standard errors are corrected for the two stage procedure, Table S4 in Appendix D shows that if the estimated ancillary parameters of the CCPs were the true parameters, both net income and physical ailments would be highly significant. We highlight this fact because the final specification in (13) is informed by the statistical significance of results before the correction of standard errors, which is computationally intensive and is done only in the end.

zero the impact of ailments for those consumers using a treatment. Therefore, the positive utility parameter for not suffering ailments while not consuming a treatment is relatively large. In other words, the cost of ailments is larger when consumers are not consuming a treatment. This finding is consistent with the idea that the utility cost of ailments from side effects of medical treatment may be less than the cost of ailments due to illness. One possible reason why is that consumers perceive ailments due to product side effects as an indicator of that effective treatments are working. Alternatively, ailments due to sickness may be more frightening than those due to medication that treats sickness.

Utility parameters for treatment choices are interacted with race and age. Since the non-pecuniary flow utility from no treatment is normalized to zero across groups, the non-pecuniary flow utility for different groups is relative to what they gain from not taking a treatment. According to parameter estimates, once we have accounted for dynamic payoffs, using treatments is costly for all consumers, with higher costs accruing to non-white consumers. Black men face a particularly high penalty of trial participation, a finding that is consistent with a broad literature investigating historical reasons why African Americans are reluctant to enter trials to use experimental drugs (Harris et al., 1996; Alsan and Wanamaker, 2016). However, age helps to mitigate the utility costs of treatment, reflecting how older agents are more accustomed to taking medications or have more contact with the medical community.

Recall that health is interacted with treatment choices that involve experimentation (either participation in a clinical trial or choosing a new drug through a cluster). We find that better health leads to larger utility costs of experimentation. This is consistent with more frequent contact with doctors among less healthy patients, who may thus face lower costs of switching to new or experimental treatments. In the case of trials, there may be more slots available for sicker patients if a goal is to test drugs on patients who most need them.

Finally, the utility of remaining on a treatment is positive, which means that it is preferred over taking no treatment at all (assuming the individual is suffering from ailments, which is normalized to zero) and over choosing a cluster or trial treatment. This underscores the idea that individuals are reluctant to experiment with new drugs. Taken together, utility estimates imply that the highest flow utility accrues to individuals who are not suffering ailments and who are not taking any treatment.

5.3 Transitions and Outcomes

Next, we discuss the processes describing how state variables produce outcomes or transition to other states. Processes for health and ailments are estimated together with product characteristics using equations (S5) and (S6) in Appendix C. Given the large number of treatments, to conserve on space, we present estimated treatment characteristics in Table S5 in Appendix D.[40] Apart from specific treatments, current health also affects current physical ailments along with future health. To capture non-linearities in these relationships, we condition on a fifth-order polynomial in health. The estimated relationships (see Table S6 in Appendix D) are plotted in the top panels of Figure 8. While the slight concavity of the production function for health could be well approximated by a linear function, the production function for ailments is very non-linear. The figures suggest that in the region of CD4 counts below 250 units, changes in health generate much larger larger shifts in the log odds ratio of suffering ailments. The reason is that HIV infection has a gradual negative impact on immune system health, as measured by CD4 count. However, the impact of CD4 count on ailments is not gradual. It is virtually non-existent until CD4 count has dropped below about 250 and AIDS-related symptoms emerge.

We next consider processes governing income, out-of-pocket payments, labor supply and survival. As before, instead of interpreting parameters on a high-order polynomial in health, we simply plot the non-linear relationships in the remaining panels of Figure 8. According to the figure, health exhibits strongly non-linear relationships with other outcomes, which helps to explain sharp differences in optimal choices for individuals with fairly similar CD4 counts. Again, this is due to large changes in physical health once the AIDS threshold is reached. These relationships underscore the importance of modeling the relationship between health and outcomes in a non-linear fashion for HIV-positive individuals.

Estimates for processes governing income, out-of-pocket payments, labor supply and survival are found in Tables 6, 7, 8 and Table 9, respectively. Beyond the relationships with health discussed above, several key patterns emerge. Individuals who do not suffer ailments have higher income as their productivity is likely to be higher. Income is concave in age and it increases with employment and education, though racial minorities earn less on average. Conditional on positive out-of-pocket medical expenditures, these payments increase with age. Minorities spend less and more educated people spend more. Similarly, individuals that

[40]Product characteristics also determined the process of within cluster assignment. As mentioned in Section 3, we do not model the shares of products within clusters. Instead, individuals who switch from a market product, a trial product, or no product, can choose a cluster of similar products and are assigned a product within the cluster, where the assignment probabilities are determined by the characteristics of the products in the cluster. Estimates of this process, described by equations (9) and (10) are presented in Table S7 in Appendix D.

suffer ailments face higher expenditures, perhaps because they are managing other health conditions. Employment increases expected payments, which may reflect different pricing schemes for public versus private insurance. Labor force participation is stochastic in our model and it is revealed to individuals at the beginning of the period. Estimates show that the log odds ratio of working versus not working increases with age until about age 40, after which point it decreases. The odds of working increase with education. Moreover, there is strong persistence in employment, reflected by a large increase in employment odds for individuals who worked in the previous period. At the end of every period individuals face the possibility of death. Estimates in imply that the log odds ratio of death versus survival decreases with age until about age 35 and then increases. The likelihood of death is smaller for black individuals and for individuals who are not suffering ailments.

5.4 Simulated Choice Dynamics and Model Fit

In Figure 9, we plot observed treatment choices over time along with those generated by the model given the state at every point in time.[41] The estimated model captures key trends remarkably well, including the rise in treatment usage as drugs improve through innovation, and trial participation dynamics. We cannot fully reproduce the spike in trial participation shortly before HAART introduction. The reason for this may be that, although our model accounts for changes in the demand for trials, there was also a shift in the supply of trials as a number of new drugs were tested that would eventually comprise the breakthrough. Hence, the spike in trial participation would not be fully captured by our model as it focuses on patient demand.

6 The Evolution of Technology

Through the lens of our model, the observed path of technological innovation is a draw from an underlying stochastic process that relates the supply of new treatments to consumer behavior. Because we have estimated this stochastic process, conditional on an aggregate set of state variables, we can simulate counterfactual technological paths and compare them to the realized path we observe. In Section 6.1), we assess the probability of the observed path by comparing it to a distribution of simulated paths using the estimated parameters of the innovation process. In Section 6.2, we examine how consumer behavior affects the distribution. In particular, we show that demand can slow the progress of innovation due to consumer reluctance to use new drugs, which affects the position of the centroid. Moreover,

[41] The fit of our parametric ccps is discussed in Appendix C.

consumer distaste for side effects can tilt the path of innovation towards new products with fewer side effects, but lower efficacy.

6.1 The Distribution of Technology Paths

We first assess the likelihood of the observed innovation path compared to what would have been predicted by the estimated innovation process. As the innovation process is used to model consumer beliefs about future innovations, this exercise also amounts to comparing the realized path of innovations to what a rational consumer would have expected starting at different points in time as the market for HIV drugs mature. To simulate technology paths, we use estimated parameters of the innovation process and draw from the non-parametric distribution F_ν. We simulate 100 innovation paths until we have reached the end of the realized path of innovations in 2008. We use two different observed periods as initial states: the first semester of 1991 and the second semester of 1996. 1991 captures the state of the market prior to the introduction of breakthroughs and health among HIV+ men was in decline. By 1996, HAART had been introduced, which led to upward shifts in average health of HIV+ men. Technology paths are plotted in Figures 10 and 11, where the solid line is the realized average and the light grey lines are the 100 simulates.

In Figure 10, we compare realized and simulated average health, ailments and survival probabilities. The panels on the left use 1991 market characteristics as starting values and the right panels start in 1996. Figure 11 shows centroid quantities and market shares over time. The upper-left panel of Figure 10 CD4 counts using 1991 as the starting point. In the first few years, simulated average health tracks realized average health quite closely, though is on average a bit worse. A larger difference emerges after 1996, however, once HAART in introduced. The reason is that HAART marks a departure from the expected path of innovations and thus average health, which the simulations show. In other words, the estimated stochastic process treats the introduction as a low probability even, though by no means treats it as a zero-probability event. Indeed, though most of the simulated paths starting in 1991 underperform relative to HAART in the health dimension, there is a small number of paths that outperform HAART. Thus, with some probability, innovations might have emerged with even higher efficacy than HAART. In we instead use 1996 as the starting point (see the panel on the top right), the market generally underperforms relative to the simulations. Innovations did not raise average health after about 2000 even though the innovation process would have predicted steady improvements to efficacy and thus higher average health by 2005.

In the second row and third rows of Figure 10, we repeat the exercise with physical

ailments and average probability of survival, respectively. The observed share of individuals suffering physical ailments is somewhat higher than the simulated paths starting in 1991. The introduction of HAART led more individuals to use medication than expected, which increased average ailments over time. If the starting point is 1996, however, the observe path of ailments is similar to the average from the simulations, which means rational agents forming expectations in 1996 would have expected population averages to evolve as they did. Simulated survival in the population is consistent with HAART being a surprise. Rational consumers forming beliefs in 1991 on average would expect lower survival than occurred due to the introduction of HAART. Starting in 1996, however, the same consumers would have expected slightly higher survival rates arising from continued innovations, which the market did not deliver, again due to the lack of subsequent innovations that improved efficacy.

Turning to market aggregates, Figure 11 simulates two quantities that determine the position of the centroid, the first capturing the weighted average of efficacy and the second capturing market average of side effects technology. Beginning in 1996, simulated paths of technology in the efficacy dimension show a trend similar to the realized path. However, HAART is a tail event and after it is introduced, the market outperforms most simulated distributions. Beginning in 1996, however, the market underperforms relative to expectations as the centroid is expected to continue moving towards higher efficacy drugs, which is not observed. For side effects innovations, the market underperforms if we simulate beginning in 1991 and 1996. An exception is that around 2001, when side effects innovation begin to occur, which means that by about 2005, the market outperforms relative to rational expectations formed in 1996. Finally, we consider market shares. Simulated market shares starting in 1991 track observed usage fairly well until about 1996. However, even without HAART, market shares rise in the years after 1997, likely because individual health would have been expected to continue its decline, which would have driven patients onto market drugs even if they were not very effective. This is in line with earlier research showing that individuals avoid effective drugs with side effects when in good health, but will use less effective drugs in spite of side effects when in poor health (Papageorge, 2016).

6.2 Demand Pull: How Consumer Choices Affect Innovation

In the structural model, the distribution of innovations is a function of aggregate consumer behavior. To illustrate this feature of the model, we consider two policies. First, we study the evolution of product quality when the process of innovation is independent of demand. Second, we assess product quality when demand is allowed to affect innovation, but consumers are randomly assigned to products. Both policies exogenously separate consumer

dynamic optimization from the process of innovation.

We start by studying innovation that is independent of consumer demand. To achieve this we redefine the centroid to be a simple average of products in the market and reestimate the process of innovation in Section 5.1.[42] Thus, product entry is no longer dependent on product demand. We also separate product exit from demand by adopting new exit rules designed to resemble the actions of a scientific authority tasked with keeping only the best products on the market. We consider two ways in which products of inferior quality are exogenously removed from the market. The first method, denoted *frontier*, removes all products from the market that are not on the technological frontier. This policy provides an upper bound for how quickly innovation can proceed according to model estimates. The second method, which captures expert intervention in a more realistic fashion, is denoted *thick frontier*. In this regime, the number of products leaving the market is given by the exit rate which we set at the baseline average, and the worst products are dropped independent of demand.[43]

Results in Figure 12 indicate that exogenous product removal from the market can lead to rapid improvements in product quality. However, more realistic exogenous interventions have limited gains in quality over demand-driven selection. The figure presents simple averages of health (right) and ailments (left) qualities of the products in the market. In the figure, the solid black line is the simulated path of innovation using model estimates, where innovation is endogenous to aggregate demand. The frontier regime paths are shown using dashed lines. According to the lines, exogenously removing all but the best products speeds innovation, leading to much much better products on both dimensions of quality, which is reflected in higher average health and fewer average suffering of ailments. In contrast, under the thick frontier policy, where only the worst products are exogenously removed from the market in each period, the path of product quality evolution is not very different from the path delivered by consumer dynamic optimization, i.e., the baseline. The reason is that consumers rarely use the worst products, so removing them exogenously has relatively little impact on the position of the centroid and, therefore, on subsequent innovations. Still, the thick frontier policy does lead to health improvements. The reason is that the worst products that are removed from the market tend to be of especially low efficacy, but are still chosen by a subset of consumers since they have few side effects.

The second policy we consider assumes that consumers are randomly assigned to prod-

[42]For simulation, we still need a path of trial participation to feed into the distribution of the number of new products and the innovation process in equation (3). We use the average path from the baseline simulations.

[43]The selection of which are the worst products in the thick frontier regime is explained in Appendix D.

ucts, but otherwise does not change the innovation process, i.e., the definition of the centroid, and the distributions of innovations and number of new products. We fix the distribution for random choices to match the unconditional shares in the first year so that consumers do not respond to shifting technology or state variables. Results are shown in Figure 13 which presents average paths, computed over 500 simulations, for consumer characteristics and the status of technology captured by the centroid. The baseline solid lines in Figure 13 are the average of the grey lines in Figures 10 and 11.

One key result from this exercise is that random assignment improves health an d survival, but also leads to more physical ailments. The reason is that individuals prefer medical treatments with fewer side effects despite the detrimental impact on their health. This tilts the path of innovation towards new products with fewer side effects, but lower efficacy. In contrast, random assignment decouples consumer preferences from innovation, tilting innovation towards higher efficacy drugs, which improves health and raises survival rates. This can be seen by looking at average health, the position of the centroid or the likelihood of survival.

In general, results in this section underscore the importance of preferences in driving not only the speed, but also the direction of technological innovation. Consumers' distaste for side effects means they are willing to sacrifice efficacy for treatments with fewer side effects even if doing so increases the likelihood of poor health in the future. This behavior reveals that consumers do not make medical decisions to maximize longevity, but instead maximize lifetime utility, which includes quality of life and a lack of physical ailments. If the innovation process were independent of consumer demand, it would be difficult to argue that there is a role for a social planner to improve consumer welfare. However, as this section has shown, private choices affect public health through innovation by slowing the development of more efficacious treatments. In the next section we ask whether individual welfare can be improved by targeting the externalities with an array of interventions.

7 Policies to Address Demand Externalities

Our estimates show that individuals prefer to avoid experimentation. They face utility costs for using market products for the first time and for participating in clinical trials to access new products. These preferences affect consumer behavior and thus the evolution of technology. Moreover, since individuals are atomistic consumers do not internalize the marginal effects of their choices on the speed and direction of innovation. As a result, a social planner could improve welfare by shifting choices.

In this section we examine interventions that assign consumers to treatment options and assess which of these policies can improve average welfare. Given the size of the individual's state space, numerically solving the problem of an unrestricted planner with full information quickly becomes intractable. We simplify the analysis by exploring the nature of the externalities using a one-period planner that is constrained to act on a subset of the individuals' information or who has a reduced set of alternatives to choose from.[44] Because the planner is constrained, the magnitude of the welfare improvement, if any, is a lower bound to what an unconstrained planner could achieve. First, we consider policies where the social planner has limited information about consumer health and previous-period treatment choices. Second, we examine a planner that has full information about observed consumer characteristics, but whose only policy tool is to shift the probability that some consumers are randomly assigned to participate in clinical trials.

7.1 Optimal Assignment Rules under Limited Information

The first constrained one-period planner we consider assigns individuals to choices on the basis of their health and their previous treatment. In particular, when assigning individuals to choices, the planner considers two levels of health (high and low) and two categories of previous treatment choices (market treatment and either no treatment or a trial treatment). Compared to the amount of information in the individual's state $\mathcal{Z}_{it}$, the amount of information available to the constrained planner is minimal. Nevertheless, adding more dimensions of information increases very rapidly the number of assignment rules that we would need to evaluate.

For every individual in a $\langle health,\ previous\ treatment \rangle$ category, the planner assigns one of the six choices available or, alternatively, allows the the individual act freely. Given the structure of the model, the planner can impose that a consumer stay on treatment only if the consumer chose a market treatment in the prior period. Hence, the constrained planner can choose one of $7^2 * 6^2 = 1,764$ possible assignment rules. Since the constrained planner has less information than the individual, it may be optimal for her to let some groups act freely even though those groups will not internalize the full costs or benefits of their behavior on the evolution of the technology. This setup is attractive because it allows the planner to choose the same solution as atomistic individuals. We solve the problem of the one-period planner in the first semester of 1991. To do this, we compute lifetime utility ten times for

[44] An additional reason behind our choice of one-period planner policies is that they are unexpected shocks to consumers. Hence, we can still compute lifetime utility using the CCPs estimated in Section 4. Otherwise, individuals would adjust their choices anticipating planner policies, rendering our estimated CCPs invalid for assessment of lifetime utility under counterfactual policies.

each consumer in 1991 for each of the 1,764 possible assignment rules.

Results from this exercise are found in Table 10, where we present details for the top ten and the bottom ten assignment rules in terms of average welfare along with the atomistic solution. According to estimates, average welfare at the atomistic solution is \$346,110. The best possible planner rule generates average welfare of \$351.990, which means that the planner can deliver a small welfare improvement of about 2% over atomistic agents who make individually rational choices. This is a lower bound for the size of the externality. The worst assignment rules can lead to average welfare that is roughly 60% lower than the atomistic solution, with the worst plan generating average welfare of \$166,230.

Table 10 also provides details about each planner solution. Recall, the planner observes enough information to categorize consumers into four categories. 50% are in good health with no previous treatment in the prior period. 26% are in good health and chose a market treatment last period. Low-health non-treatment users and treatment users in the previous period comprise 5% and 19% of patients, respectively. In the best rule, individuals in good health with no prior treatment information are assigned to no treatment. Those in good health who chose a market treatment in the prior period are left to freely choose. Notice, consumers who like the qualities of the treatments they have just chosen are able to choose the same products, while those who did not like the treatment can choose something else. This is an important component of the assignment rule since the planner is essentially exploiting the fact that these individuals have information about cluster treatments that the planner does not. The third group of individuals, those in poor health and with no prior treatment experience, are likewise left to choose freely in the optimal plan. Finally, individuals in poor health who chose a market treatment in the prior period are assigned the no treatment option.

Many of the top ten rules exhibit similar features, with some exceptions. For example, in some cases, the planner assigns individuals with prior experience to their previous choice, thus avoiding the costs of experimentation. Moreover, a few of the top ten rules involve experimentation, though this only occurs among the relatively small group of sicker patients, for whom costs of experimentation are relatively low. Many of the bottom ten assignment rules share two common features. One, the planner imposes experimentation onto healthy patients who dislike experimentation the most. Two, the planner discards the information possessed by individuals who used treatments in the prior period, by reassigning these individuals to new treatments rather than assigning them to their previous treatment (thus avoiding experimentation costs) or allowing them to choose freely.

7.2 Optimal Trial Participation

In our second exercise we focus on the externality that arises from atomistic individuals who do not incorporate the marginal effects of their clinical experimentation on the evolution of technology, and who are then likely to participate in clinical trials in levels that do not maximize average welfare. One reason we focus on trial participation is that, despite the social benefits of clinical trials, individuals are generally not allowed to be financially compensated for participation. Since the market fails to price an activity that generates social benefits, an inefficient level of experimentation can arise.

We consider a one-time, constrained planner who observes the individuals' states but whose only available policy tool is to assign individuals to clinical trials or to let them choose freely among the other alternatives. To assign individuals optimally, the planner ranks individuals according to how much lifetime utility they loose if they are assigned to a clinical trial. At the top of the ranking are individuals who would have chosen a trial in the absence of the planner. These individuals loose nothing and are ranked based on their gains over their second best alternative. Next are individuals who loose welfare from being assigned to a trial. They are ranked based on their losses over their best alternative. When increasing trial participation from the atomistic equilibrium, the planner begins with individuals whose welfare losses are small. This problem nests the atomistic solution if the planner assigns trial participation solely to individuals who would have chosen to participate on their own. We solve the constrained planner's problem in the first semester of 1991 and again in the second semester of 1996.[45]

Results are presented in Table 11. In 1991, the planner's optimal share of participation in clinical trials is virtually identical to the atomistic share, yielding an average welfare of about \$345,000. This result suggests that the costs of increased experimentation outweigh the benefits of new drugs in a time when no good treatments have been invented and previous innovations have been small. Given the location of the centroid and the maximum size of previous innovations, assigning individuals to trials does not improve welfare sufficiently to justify the costs of increased experimentation.

By 1996, large innovations had occurred and further innovations were therefore more probable. Conditional on the state of the market at the second semester of 1996, the planner's optimal share is twice as large as the atomistic share and provides an average welfare of about \$360,000, about 2% higher than the atomistic average welfare. The solid line in Figure 14 indicates average welfare for different assignment rates above the atomistic share in 1996,

[45]To find the optimal rate of trial participation, for each rate, we simulate the evolution of the market 1000 times and average per person lifetime utility over all simulations.

showing that the maximum is roughly 19%, compared to the atomistic rate of 9%. At rates between 9% and 19% the gain from increased participation outweighs individual losses due to experimentation. However, at rates above 19%, average welfare drops precipitously. This is because individuals are assigned to trials who face larger losses relative to their optimal choice, and the returns to additional experimentation are not large enough.

To measure the magnitude of the externality in 1996, we compute the net social benefit from assigning the marginal consumer, to a trial. This is the consumer who would otherwise choose something else, but who faces the smallest lifetime utility loss from assignment to a trial. This is equivalent to computing the derivative average welfare with respect to the trials share, evaluated at the atomistic share. We find that the marginal consumer loses roughly \$600 (Table 11). However, because trial participation spurs innovation by pushing up the expected quality and the expected number of new products, the net social gain is over \$2,000 per person. In our sample of 445 individuals in 1996, this means that a \$600 loss from raising trial participation by 1 person (roughly 0.22 percentage points) leads to a welfare gain of roughly \$1,000,000.[46]

These results suggest a substantial externality. However, it is doubtful that a government agency would attain the necessary information to rank individuals by unobservable utility and thus be able to minimize losses. Moreover, there are ethical concerns if a government body ever attains the authority to assign individuals to clinical experimentation. As history has shown, who is sent to experimentation depends heavily on the weights allocated to individuals of a given group in the social welfare function (Harris et al., 1996).

A more realistic and ethical solution would be to provide economic incentives for individuals to participate in clinical trials. We study a flat subsidy that is paid to all participants in clinical trials, including individuals who would have chosen to join a trial absent a subsidy. When using a flat subsidy to induce more experimentation in clinical trials, the government agency faces an increasing cost per participant. The reason is that the amount needed to attract the highest-cost participant must be paid to all enrollees. Figure 14 shows that the optimal share attained through a flat subsidy is 13.5%, which is substantially lower than the optimal trial share of the informed planner of 18.5% (Table 11). The subsidy needed to induce this level of experimentation is \$8,500 per trial participant. Despite these costs, the flat subsidy increases average lifetime utility (net of the subsidy) to \$357,000 which is about 1% above the atomistic average value of \$354,000 or about \$3,000 per person. Therefore, our results suggests that providing monetary incentives for trial participation can be welfare improving.

[46]Since the line in Figure 14 is fairly bumpy, we use a fifth degree local smoothing polynomial to evaluate the marginal gains.

8 Conclusion

We build a structural model to assess how consumer choices affect the evolution of technology in the market for HIV drugs. We capture several mechanisms through which consumer demand affects innovation, including experimentation with new drugs by participating in clinical trials, which accelerates innovation. By joining a trial, individuals gain access to experimental products that may be high-quality breakthroughs, but may also be less efficacious or painfully toxic. Additionally, consumer decisions can bend the technological path if firms avoid innovating around unpopular products. Because individuals are atomistic, an externality arises.

Our results show that consumer behavior can slow the process of innovation and bend it towards less efficacious products that hamper survival probabilities. They also show that atomistic individuals sometimes engage in levels of experimentation that are not welfare maximizing. We explored these issues by considering constrained planner problems, which provides lower bound measures for the size of the externality. We find that a constrained planner can increase average welfare by around two percent (approximately $6,000 per individual). Additionally, we show that welfare-enhancing levels of trial participation can be achieved using a flat subsidy. In other words, our results suggest that providing monetary incentives for trial participation can be welfare improving.

Given our data, we have focused on how demand and consumer heterogeneity affect the path of technological progress. Other studies have placed more emphasis on the supply side (Carranza, 2010; Goettler and Gordon, 2011; Gowrisankaran and Rysman, 2012). A natural step forward, although by no means a simple one, is to model competition among firms and firm decisions to innovate by investing in R&D, while also maintaining an acceptable level of consumer heterogeneity and a role for demand pull. This approach would allow for a richer set of counterfactual policies that incorporate the interaction between consumer demand and firm strategic behavior.

References

Allcott, Hunt, Rebecca Diamon, and Jean-Pierre Dubé, "The Geography of Poverty and Nutrition: Food Deserts and Food Choices Across the United States," Mimeo, New York University 2015.

Alsan, Marcella and Marianne Wanamaker, "Tuskegee and the Health of Black Men," 2016.

Altuğ, Sumru and Robert A. Miller, "The Effect of Work Experience on Female Wages and Labour Supply," *Review of Economic Studies*, 1998, *65*, 45–85.

Arcidiacono, Peter and Robert A. Miller, "Identifying Dynamic Discrete Choice Models off Short Panels," Working Paper 2015.

Bayer, Patrick and Robert McMillan, "Tiebout Sorting and Neighborhood Stratification," *Journal of Public Economics*, 2012, *96* (11), 1129–1143.

Bhaskaran, K, O Hamouda, M Sannes, and et al, "Changes in the Risk of Death after HIV Seroconversion Compared with Mortality in the General Population," *Journal of the American Medical Association*, 2008, *300* (1), 51–59.

Bhattacharya, Jay and Mikko Packalen, "The Other Ex Ante Moral Hazard in Health," *Journal of Health Economics*, 2012, *31* (1), 135–146.

Bolton, Patrick and Christopher Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Carranza, Juan Esteban, "Product Innovation and Adoption in Market Equilibrium: The Case of Digital Cameras," *International Journal of Industrial Organization*, 2010, *28*, 604–618.

Chan, T.Y. and B.H. Hamilton, "Learning, Private Information, and the Economic Evaluation of Randomized Experiments," *Journal of Political Economy*, 2006, *114* (6), 997–1040.

Crawford, G.S. and M. Shum, "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 2005, pp. 1137–1173.

Dranove, David, Craig Garthwaite, and Manuel Hermosilla, "Pharmaceutical Profits and the Social Value of Innovation," NBER Working Paper 2014.

Duda, R.O. and P.E. Hart, *Pattern Classification and Scene Analysis*, Wiley, 1973.

Erdem, Tülin and Michael P Keane, "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 1996, *15* (1), 1–20.

Finkelstein, Amy, "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry," *Quarterly Journal of Economics*, 2004, *119* (2), 527–564.

Goettler, Ronald L. and Brett Gordon, "Does AMD Spur Intel to innovate More?," *Journal of Political Economy*, 2011, *119* (6), 1141–1200.

Gowrisankaran, Gautam and Marc Rysman, "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy*, 2012, *120* (6), 1173–1219.

Harris, Yvonne, Philip B Gorelick, Patricia Samuels, and Isaac Bempong, "Why African Americans May Not Be Participating in Clinical Trials.," *Journal of the National Medical Association*, 1996, *88* (10), 630.

Hicks, John R, *The Theory of Wages*, London: Macmillan, 1932.

Hincapié, Andrés, "Entrepreneurship over the Life Cycle: Where are the Young Entrepreneurs?," *Mimeo, University of North Carolina at Chapel Hill*, 2017.

Hotz, V Joseph and Robert A Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

_ , _ , Seth Sanders, and Jeffrey Smith, "A Simulation Estimator for Dynamic Models of Discrete Choice," *The Review of Economic Studies*, 1994, *61* (2), 265–289.

Ickovics, Jeannette R and Andrew W Meisler, "Adherence in AIDS Clinical Trials: A Framework for Clinical Research and Clinical Care," *Journal of Clinical Epidemiology*, 1997, *50* (4), 385–391.

Jovanovic, Boyan and Glenn M. MacDonald, "Competitive Diffusion," *Journal of Political Economy*, 1994, *102* (1).

Lancaster, K.J., "A New Approach to Consumer Theory," *The Journal of Political Economy*, 1966, *74* (2), 132–157.

Magnac, Thierry and David Thesmar, "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 2002, *70* (2), 801–816.

Miller, Robert A., "Job Matching and Occupational Choice," *The Journal of Political Economy*, 1984, *92* (6), 1086–1120.

_ , "Innovation and Reputation," *Journal of Political Economy*, 1988, *96* (4), 741–765.

Papageorge, Nicholas W., "Why Medical Innovation is Valuable: Health, Human Capital, and the Labor Market," *Quantitative Economics*, 2016, *7* (3), 671–725.

Petrin, A., "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 2002, *110* (4), 705–729.

Romer, Paul M, "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 1986, pp. 1002–1037.

Rosen, S., "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *The Journal of Political Economy*, 1974, *82* (1), 34–55.

Scherer, Frederic M, "Demand-Pull and Technological Invention: Schmookler Revisted," *The Journal of Industrial Economics*, 1982, pp. 225–237.

Schmookler, Jacob, *Invention and Economic Growth* 1966.

Stigler, G.J., "The Cost of Subsistence," *Journal of Farm Economics*, 1945, *27* (2), 303–314.

Waldfogel, Joel, "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *RAND Journal of Economics*, 2003, pp. 557–568.

9 Figures and Tables

Table 1: Summary Statistics: Subjects. Visit 14-47 (1990-2007)

	Restricted Sample	
Subjects	1719	
	mean	std dev
Black	0.22	
Hispanic	0.09	
White	0.68	
High School	0.14	
Some College	0.29	
College	0.34	
More than College	0.23	
Died	0.40	
Died Conditional	0.20	
Ever Take Market Product	0.83	
Ever Take Trial Product	0.24	
Ever Work	0.74	
Ever Not Work	0.68	
Age in 1991	36.04	(8.72)

Notes: Standard deviation in parentheses. Data for unique individuals. *Ever Market Product* stands for ever consumed a market product during the period from visit 14 to visit 47. Similar definition holds for *Ever Trial Product*. *Died Conditional* is the proportion of individuals who died conditional on surviving until year 1995.

Table 2: Summary Statistics: Subjects-Visits. Visits 14-47 (1990-2007)

	Analytic Sample	Pre Haart	Post Haart
Obs	16851	6972	9879
Ailments	0.43	0.45	0.41
Market Product	0.65	0.49	0.76
Trial Product	0.07	0.09	0.05
Work	0.63	0.70	0.58
Age	44.48	40.89	47.01
	(8.03)	(6.99)	(7.75)
CD4	475	407	524
	(297)	(298)	(287)
Gross Income	17567	19036	16531
	(8787)	(8733)	(8677)
Out-of-pocket Pay	266	179	327
	(706)	(598)	(767)

Notes: Standard deviation in parentheses. Income and Out-of-pocket are semestral and measured in real dollars of 2000. Pre HAART era corresponds to visit <= 24 or (roughly) year <= 1995.

Table 3: Innovation Components

Conditional Mean of Innovations ν^*

	Health Innovation ν^{*h}_{rt+1}		*Ailments Innovation* ν^{*x}_{rt+1}	
	coef.	se	coef.	se
$TrialsShare_t$	433.11	(19.95)	1.93	(0.34)
$Constant$	-24.14	(1.47)	-0.15	(0.03)

Covariance Matrix of the Distribution of Innovations F_ν

	$Health$		$NoAilments$	
	coef.	se	coef.	se
$Health$	396.07	(27.85)		
$NoAilments$	1.77	(0.34)	0.14	(0.01)

Notes: Top panel corresponds to estimates from equation (3). Bottom panel shows the covariance matrix of the distribution of innovations. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 4: Distribution of Number of New Products, F_N

	coef.	se
μ		
Q_t	0.432	(0.074)
$TrialsShare_{t-1}$	6.177	(0.495)
$\ln \alpha$		
$Constant$	-0.206	(0.051)
Q_t	-1.019	(0.139)

Notes: Model specified in (5). The variable Q_t measures the distance between the previous period's new products and the previous period's centroid. It captures the relatively higher number of new products that follow the appearance of better innovations. The variable $TrialsShare_{t-1}$ is the share of individuals going into a trial the previous period. According to the model in (5), $E[New_{t+1}] = \mu_t$ and $Var[New_{t+1}] = \mu_t(1 + \alpha\mu_t)$. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 5: Utility Parameters, y_{it}

parameter	variable	coef.	se
α_{4w}	$Cluster_{it} \cdot White_i$	-3.546	(0.744)
α_{4b}	$Cluster_{it} \cdot Black_i$	-4.190	(0.762)
α_{4l}	$Cluster_{it} \cdot Hispanic_i$	-3.967	(0.958)
α_{4a}	$Cluster_{it} \cdot Age_{it-1}$	0.043	(0.011)
α_{4h}	$Cluster_{it} \cdot Health_{it-1}/10^3$	-2.021	(0.423)
α_{5w}	$Trial_{it} \cdot White_i$	-1.468	(0.280)
α_{5b}	$Trial_{it} \cdot Black_i$	-2.553	(0.334)
α_{5l}	$Trial_{it} \cdot Hispanic_i$	-1.585	(0.356)
α_{5a}	$Trial_{it} \cdot Age_{it-1}$	0.032	(0.005)
α_{5h}	$Trial_{it} \cdot Health_{it-1}/10^3$	-2.461	(0.203)
α_{6w}	$Stay_{it} \cdot White_i$	0.502	(0.567)
α_{6b}	$Stay_{it} \cdot Black_i$	0.276	(0.613)
α_{6l}	$Stay_{it} \cdot Hispanic_i$	0.707	(0.454)
α_{6a}	$Stay_{it} \cdot Age_{it-1}$	0.009	(0.007)
α_{xp}	$NoAilments_{it} \cdot NoProduct_{it}$	1.019	(1.767)
α_m	$GrossIncome_{it} - OutPocketPay_{it}$	0.057	(0.057)

Notes: Estimation of equation (13). Discount factor $\beta = .95$. $Cluster_{it}$ indicates whether the individual chose one of the three clusters of products available. $Trial_{it}$ indicates whether he chose a trial treatment. $Stay_{it}$ indicates whether he decided to continue using the same treatment he used last period. $NoProduct_{it}$ indicates whether he did not consume a product. $Health_{it-1}$ is defined as the number of white blood cells per cubic millimeter of blood. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 6: Gross Income, m_{it}

variable	coef.	se
h_{it-1}	0.018	(0.001)
$h^2_{it-1}/10^3$	-0.064	(0.007)
$h^3_{it-1}/10^7$	1.138	(0.171)
$h^4_{it-1}/10^{10}$	-1.030	(0.213)
$h^5_{it-1}/10^{14}$	4.854	(1.414)
$h^6_{it-1}/10^{18}$	-11.270	(4.712)
$h^7_{it-1}/10^{20}$	0.101	(0.062)
a_{it-1}	0.482	(0.034)
a^2_{it-1}	-0.006	(0.0004)
black	-5.534	(0.115)
hispanic	-4.167	(0.222)
some college	2.497	(0.141)
college	5.812	(0.157)
more than college	8.203	(0.151)
l_{it}	5.738	(0.074)
x_{it}	0.207	(0.024)
constant	-2.095	(0.801)

Notes: Estimation of equation (15). Random effects regression of gross-income on covariates. m_{it} is measured in thousands of real dollars of 2000. Health is given by the CD4 count measured in hundreds of cells per microliter. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 7: Tobit Model for Out-of-pocket Payments, o_{it}

variable	coef.	se
h_{it-1}	-0.002	(0.0004)
$h_{it-1}^2/10^3$	0.009	(0.002)
$h_{it-1}^3/10^7$	-0.133	(0.032)
$h_{it-1}^4/10^{10}$	0.090	(0.029)
$h_{it-1}^5/10^{14}$	-0.266	(0.118)
$h_{it-1}^6/10^{18}$	0.279	(0.181)
a_{it-1}	0.037	(0.004)
a_{it-1}^2	-0.0002	(0.0001)
black	-0.240	(0.014)
hispanic	-0.119	(0.016)
some college	0.169	(0.016)
college	0.318	(0.018)
more than college	0.336	(0.018)
market product	0.429	(0.016)
trial product	0.313	(0.021)
l_{it-1}	0.105	(0.009)
x_{it}	-0.122	(0.008)
constant	-1.459	(0.099)
σ^o	0.862	(0.027)

Notes: Estimation of equation (16). $MarketProduct_{it} = \sum_{k=1}^{4} d_{kit}$. Out-of-pocket Payments o_{it} are measured on thousands of real dollars of 2000. Health is given by the CD4 count measured in hundreds of cells per microliter. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 8: Logit Model for Labor Supply, l_{it}

variable	coef.	se
h_{it-1}	0.009	(0.0003)
$h_{it-1}^2/10^3$	-0.013	(0.001)
$h_{it-1}^3/10^7$	0.075	(0.005)
$h_{it-1}^4/10^{10}$	-0.013	(0.002)
a_{it-1}	0.102	(0.009)
a_{it-1}^2	-0.001	(0.0001)
black	-0.168	(0.025)
hispanic	-0.040	(0.044)
some college	0.312	(0.031)
college	0.537	(0.029)
more than college	0.613	(0.033)
l_{it-1}	4.458	(0.028)
constant	-5.914	(0.190)

Notes: Estimation of equation (17). Health is given by the CD4 count measured in hundreds of cells per microliter. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 9: Logit model for Death, $1 - S_{it+1}$

variable	coef.	se
h_{it-1}	-0.028	(0.001)
$h^2_{it-1}/10^3$	0.079	(0.005)
$h^3_{it-1}/10^7$	-1.104	(0.102)
$h^4_{it-1}/10^{10}$	0.704	(0.088)
$h^5_{it-1}/10^{14}$	-1.610	(0.285)
a_{it-1}	-0.116	(0.021)
a^2_{it-1}	0.002	(0.0002)
black	-0.509	(0.069)
hispanic	0.034	(0.076)
some college	0.060	(0.057)
college	-0.353	(0.053)
more than college	-0.512	(0.060)
x_{it}	-1.140	(0.050)
constant	1.682	(0.474)

Notes: Estimation of equation (21). Health is given by the CD4 count measured in hundreds of cells per microliter. In parentheses, standard errors computed using subsampling with 100 subsamples.

Table 10: Constrained Planner: Assignment Rules

		Groups			
	Group Share	0.50	0.26	0.05	0.19
	Average Welfare ($1000)	highH, nop/trial	highH, mk	lowH, nop/trial	lowH, mk
Top ten rules	351.99	6	7	7	6
	351.41	6	7	6	6
	350.89	6	4	7	6
	350.88	6	4	6	6
	349.42	6	7	6	7
	349.31	6	7	3	6
	349.11	6	7	1	6
	348.92	6	4	5	6
	348.79	7	4	5	7
	348.65	6	7	6	4
	⋮				
Atomistic	346.11	7	7	7	7
	⋮				
Bottom ten rules	169.72	1	5	5	5
	169.32	3	5	6	5
	169.17	3	5	7	5
	168.96	3	5	1	5
	168.11	1	5	6	5
	167.56	1	5	2	5
	167.55	3	5	5	5
	167.37	1	5	1	6
	167.19	1	5	3	6
	166.23	1	5	3	5

Notes: Best and worst assignment rules of a one-time constrained planner that solves her assignment problem in 1991. Groups are determined by health status (high or low) and previous treatment status (consumed a market treatment or not). The population shares of each of the groups are shown on top of their labels. Numbers 1 to 3 correspond to the three clusters available in 1991. Number 4 corresponds to staying in previous market treatment. Number 5 stands for trial and 6 stands for no treatment. Finally, number 7 stands for individually optimal choice; in other words, the planner renounces to her right to impose a choice and lets the individual in the group decide based on their richer information set.

Table 11: Constrained Planner: Optimal Experimentation

Planner optimal trial share	0.100	0.185
Atomistic trial share	0.102	0.092
Average lifetime utility at planner solution	345	360
Average lifetime utility at atomistic equilibrium	345	354
Increment in trial share for marginal person sent to trials at atomistic equilibrium	0.001	0.002
Individual loss for marginal person sent to trials at atomistic equilibrium	-0.178	-0.628
Social gain from sending marginal person to trials at atomistic equilibrium	-1133	1051
Cost to attain marginal increment (at atomistic equilibrium) using flat subsidy	-	26.4
Flat subsidy optimal share	-	0.135
Subsidy per trials participant	-	8.50
Average lifetime utility at flat subsidy optimal share	-	357

Notes: Constrained planner's problem solved at the first semester of 1991 and the second semester of 1996. Monetary values in $1,000s.

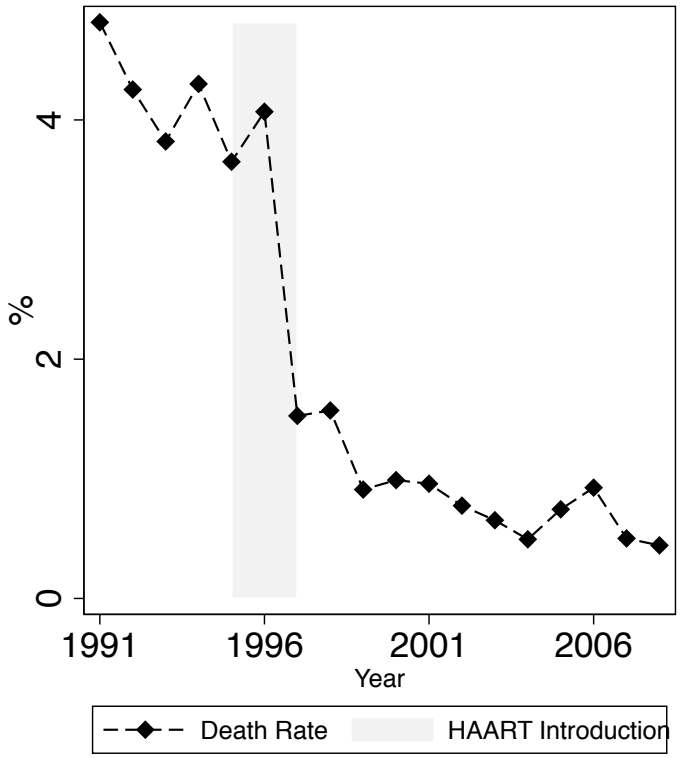

Figure 1: Death rate in the sample. More than 1500 surveyed individuals died for AIDS-related causes during our analysis period.

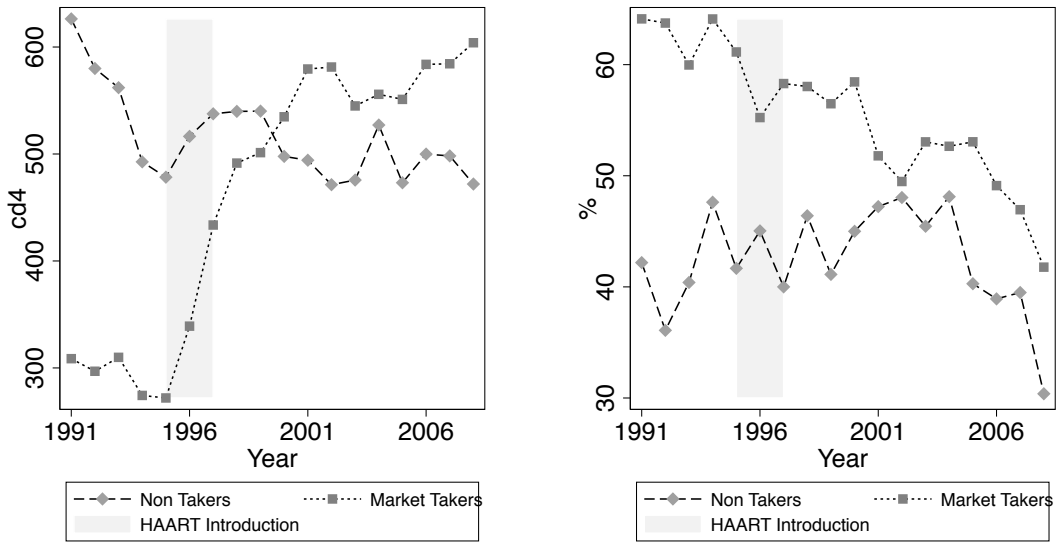

(a) (b)

Figure 2: Health and side effects summary trends over time. Panel 2(a) shows the mean CD4 over time by consumption status. Panel 2(b) contains mean ailments over time by consumption status.

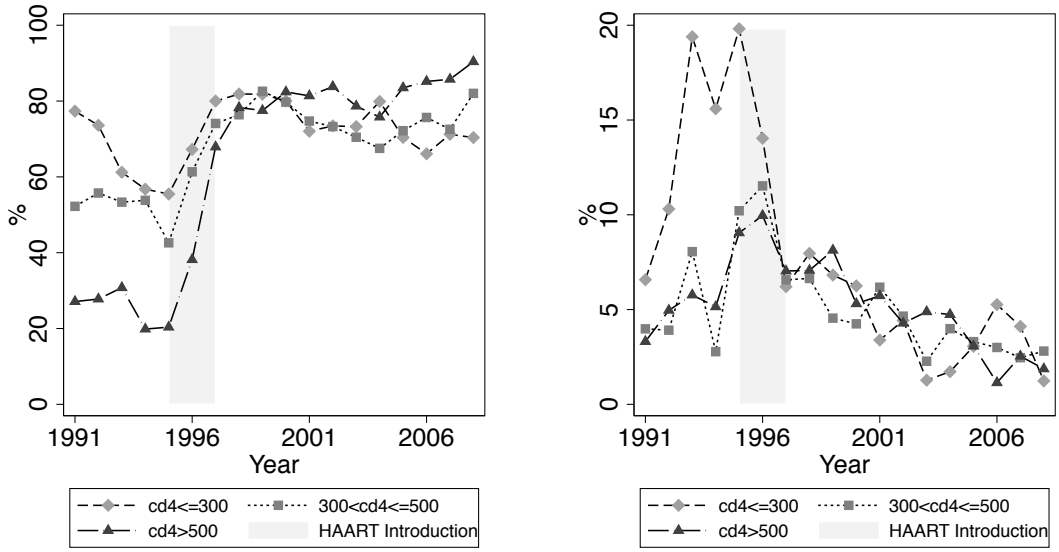

(a) (b)

Figure 3: Consumer demand over time. Panel 3(a) shows treatment consumption over time by health status. Panel 3(b) shows trial participation over time by health status.

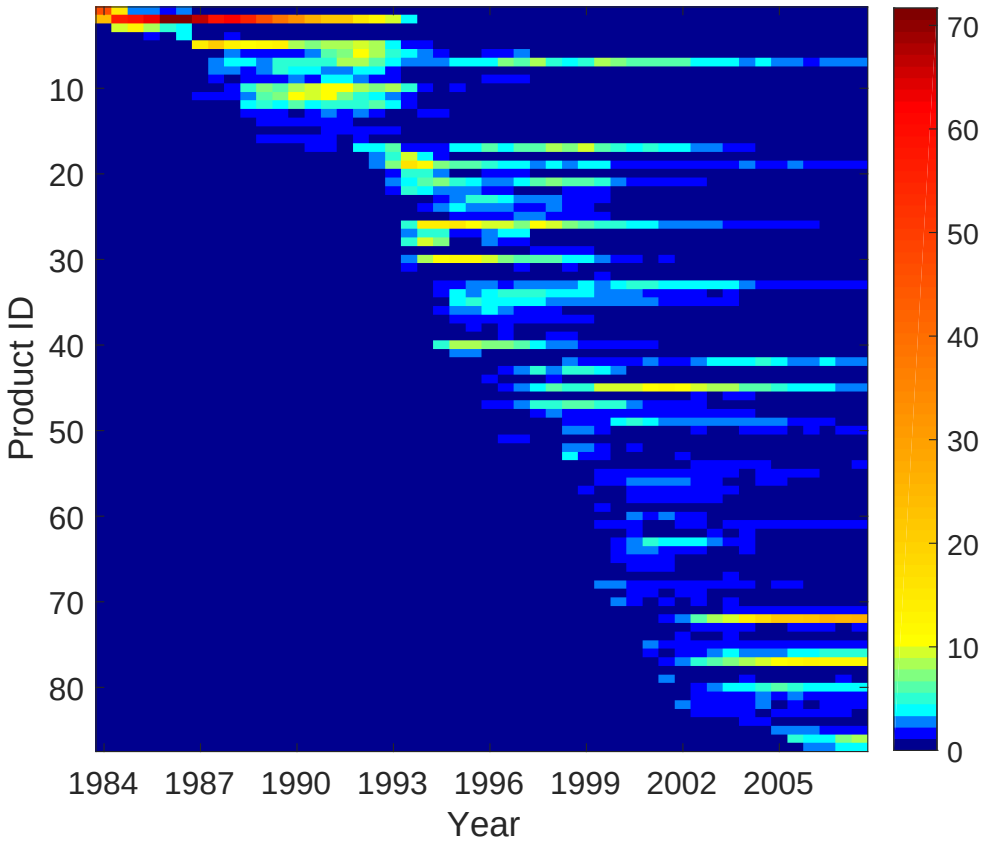

Figure 4: Diffusion of Products Over Time

Notes: HIV treatments from 1984 to 2008. Each id—or row—represents a product. Color indicates the share of the market that the product captures. Shares are conditional on consuming a product. Early on there are few products with high shares, as time passes new products strip market share from incumbents and less popular products exit.

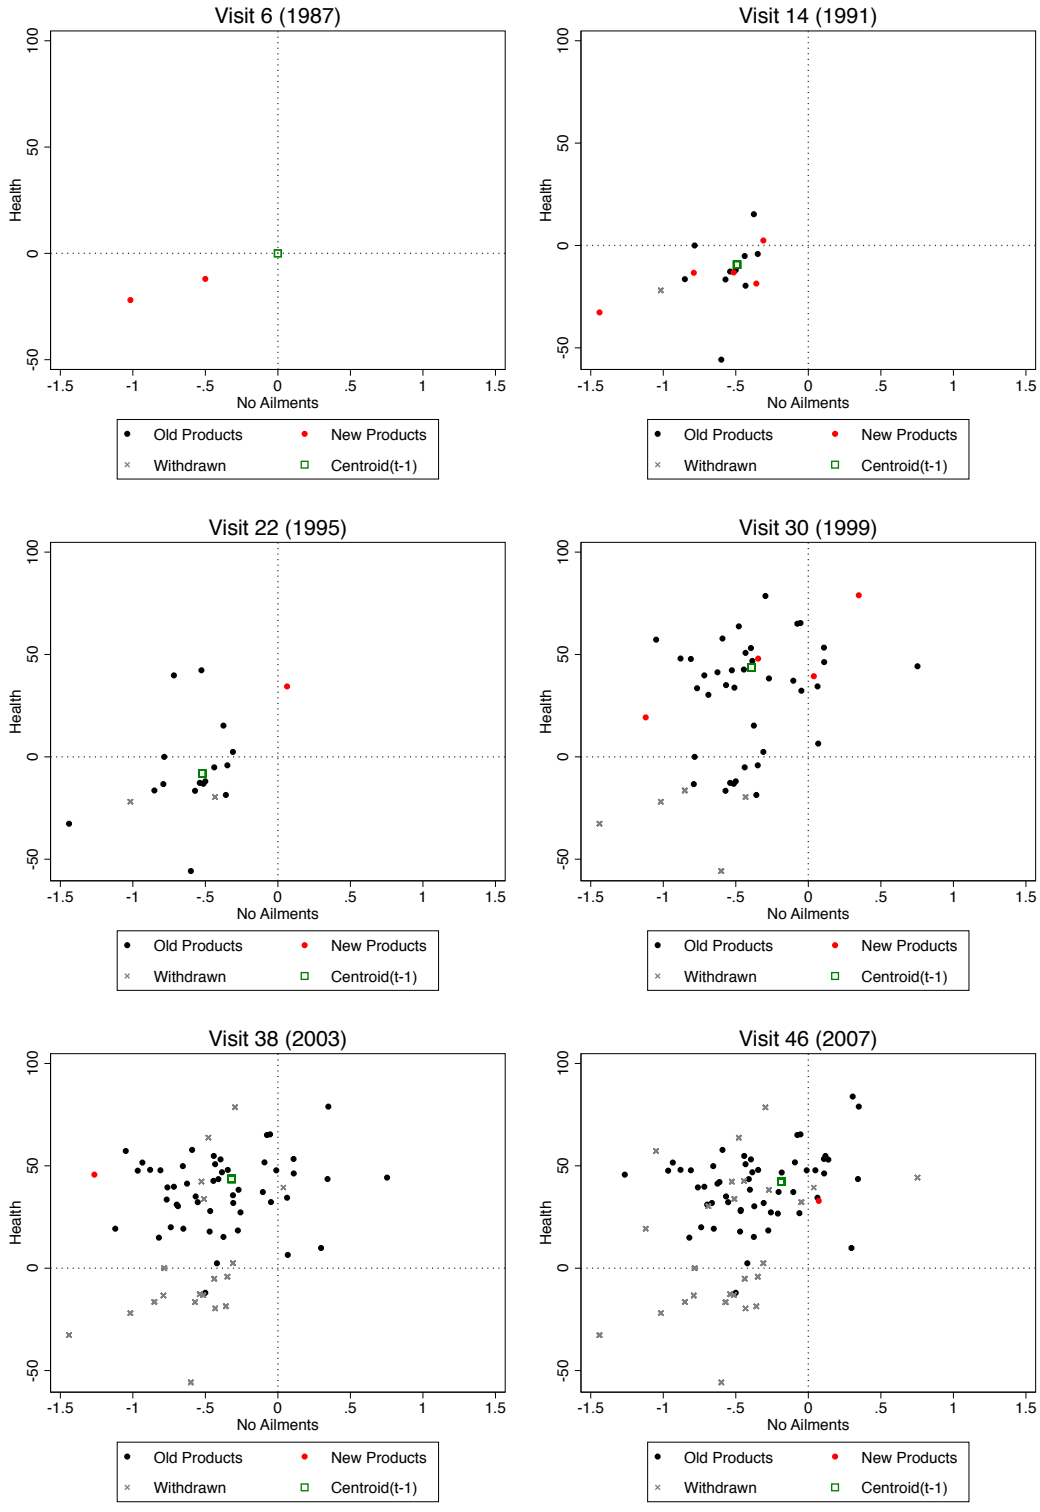

Figure 5: Treatment Evolution

Notes: Figure shows snapshots of the evolution of the state of the product market at the different stages. Products are two-dimensional. On the x-axis is a measure of a treatments ability to not cause side effects. On the y-axis is a measure of its contribution to underlying health. Dimensions are measured in different scales. Incumbent products are shown in black. New products are shown in red. Withdrawn products are shown as x. The green square is a measure of the prevalent technology in the previous period.

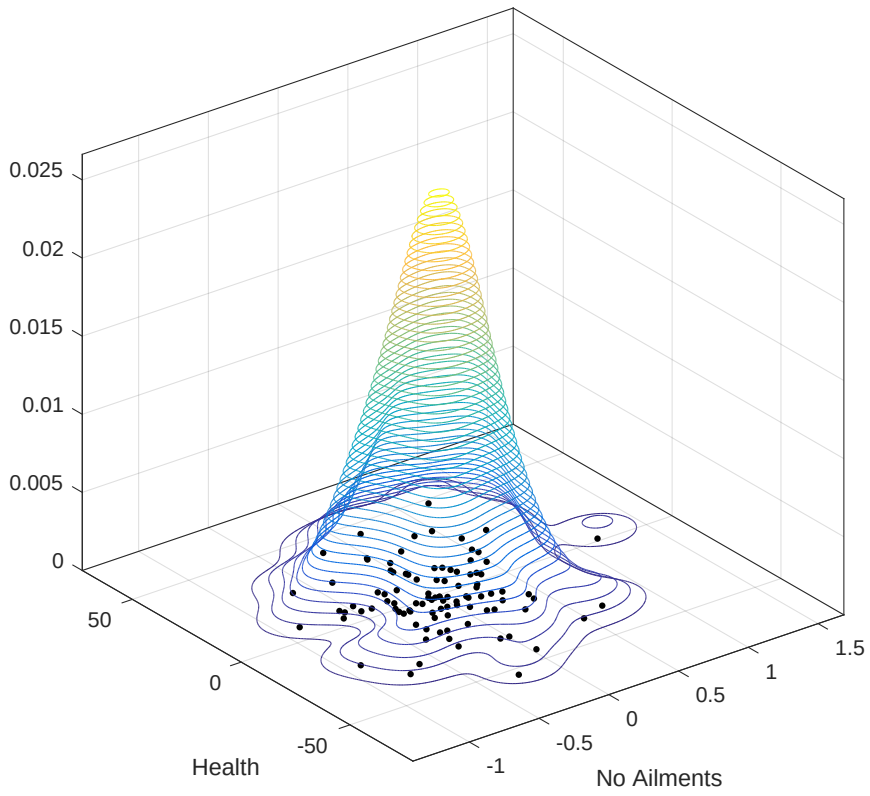

Figure 6: The Distribution of Innovations, F_ν.

Notes: F_ν is estimated non-parametrically from the realized innovation vectors ν_k implied by equations (2) and (3).

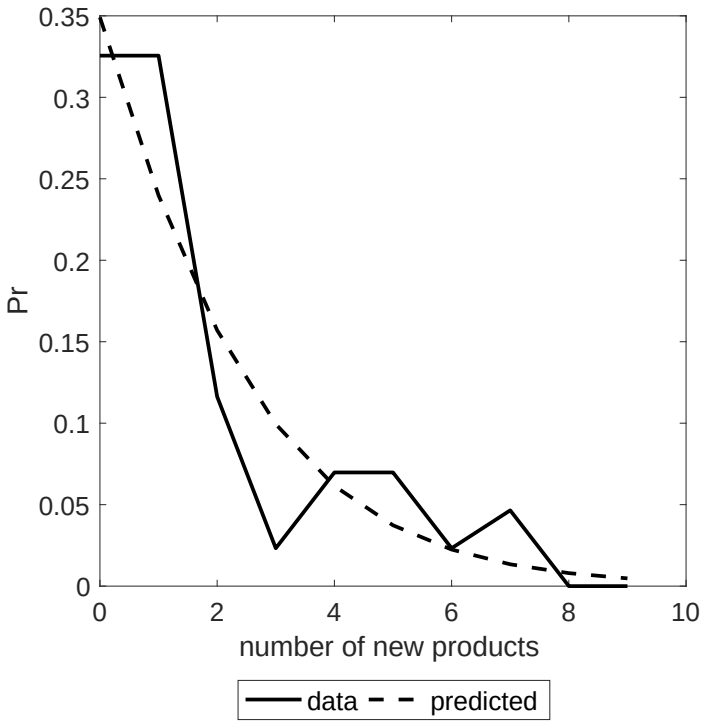

Figure 7: Distribution of Number of New Products

Notes: Model specified in (5). Figure shows the empirical distribution of new products and the average over time of the predicted probabilities using the estimated parameters in Table 4.

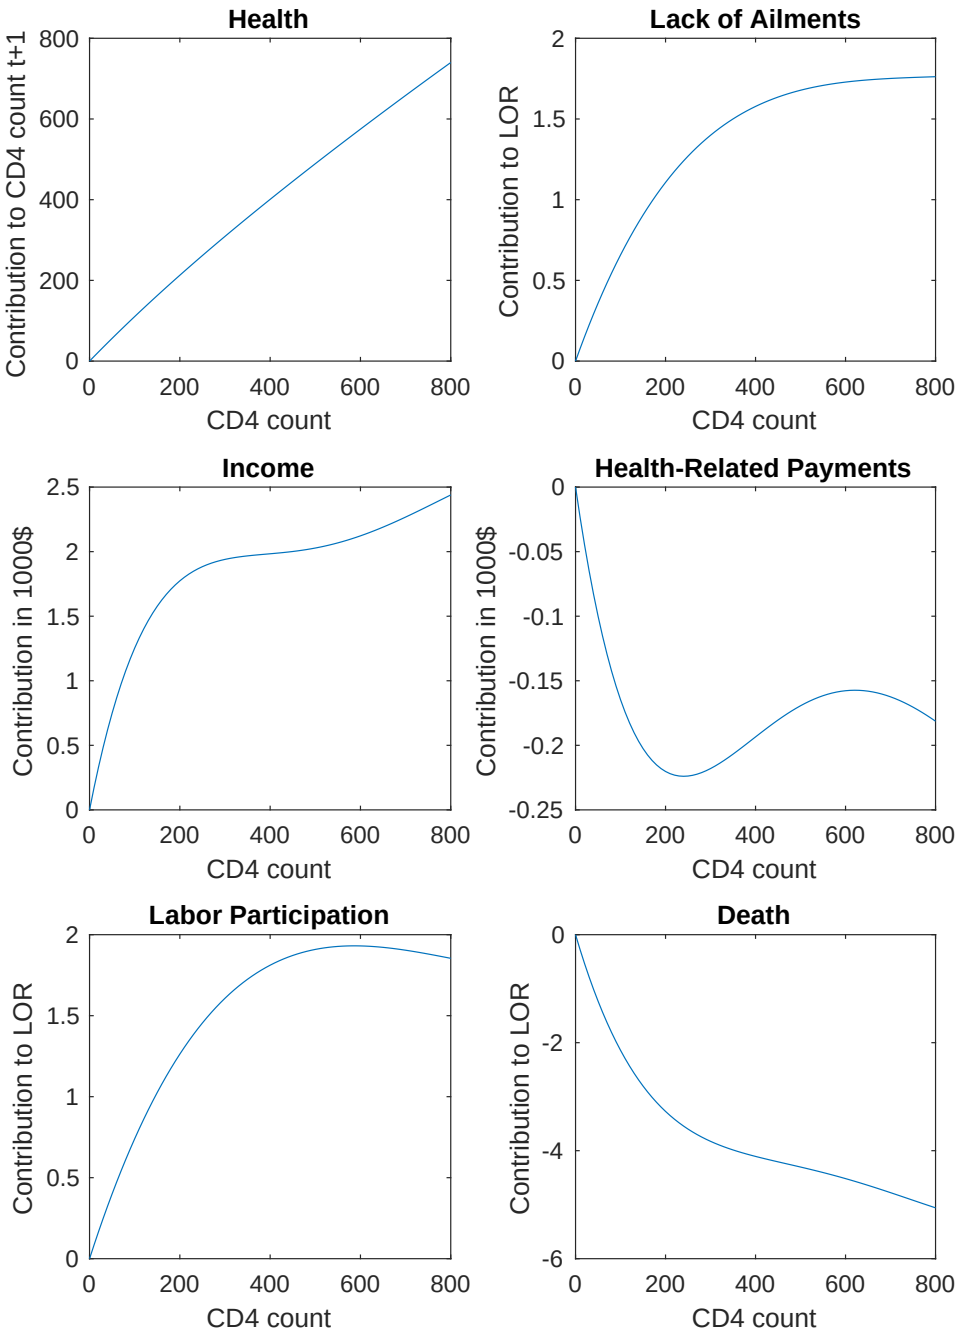

Figure 8: Health Effect on Future Health and Outcomes

Notes: CD4 Count measured in hundreds of cells per microliter. LOR stands for log odds ratio. Semestral income measured in thousands of dollars of 2000.

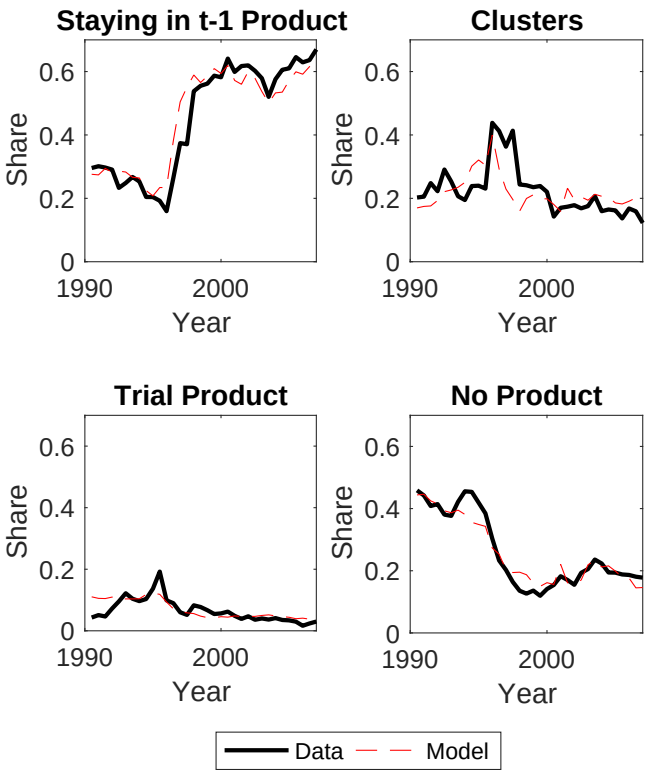

Figure 9: Goodness of Fit Figures

Notes: Simulated and empirical choice rates over time.

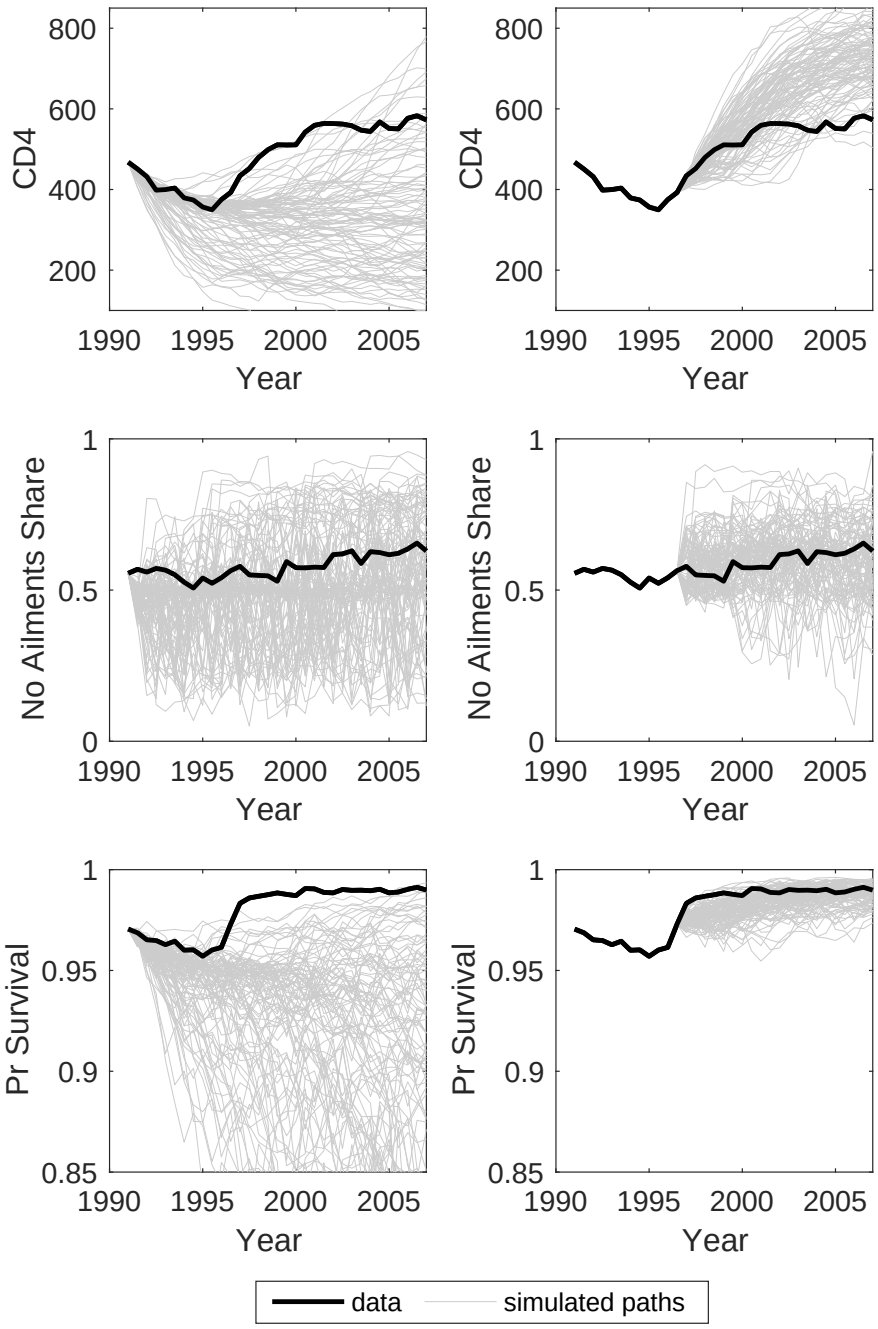

Figure 10: Distribution of Technology Paths: Consumers

Notes: 100 simulated paths conditional on the state of the world at 1991 and 1996.

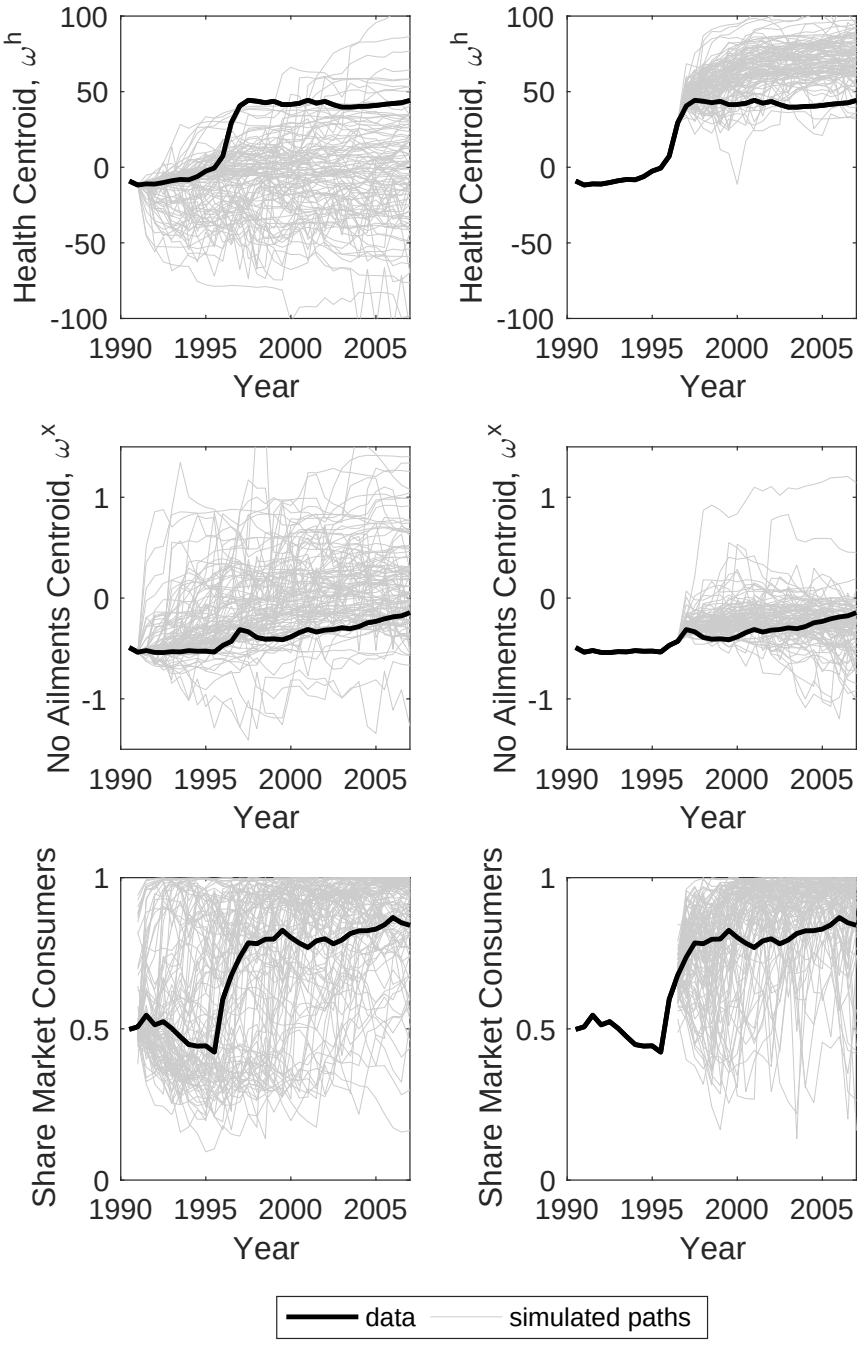

Figure 11: Distribution of Technology Paths: Technology and Product Consumption

Notes: 100 simulated paths conditional on the state of the world at 1991 and 1996.

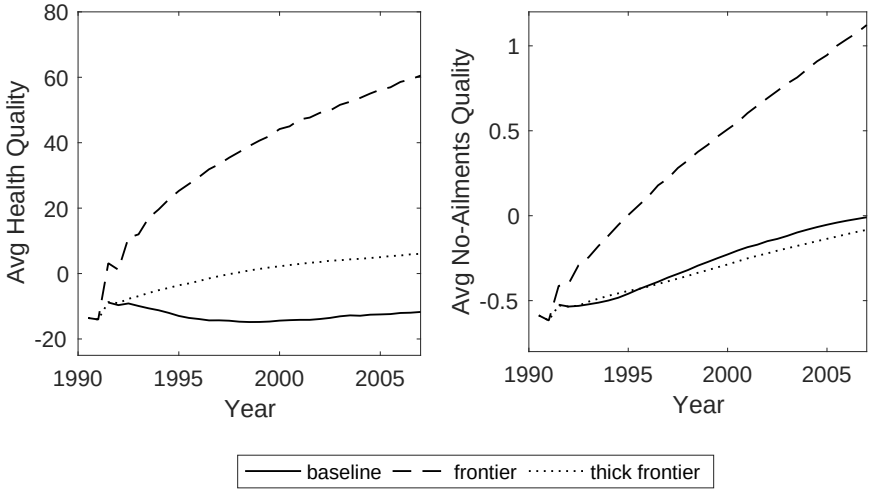

Figure 12: Alternative Regimes: Exogenous Scientific Intervention

Notes: Figure shows differences across regimes in terms of the average quality of products in the market. The *baseline* is the estimated model where products are dropped in response to low demand. In the *frontier* regime all dominated products in terms of quality are exogenously dropped from the market every period. The *thick frontier* sets the exit rate of products at the average baseline rate and exogenously drops the worst products. Statistics are computed from 500 simulations that are conditional on the state of the world at the first semester of 1991.

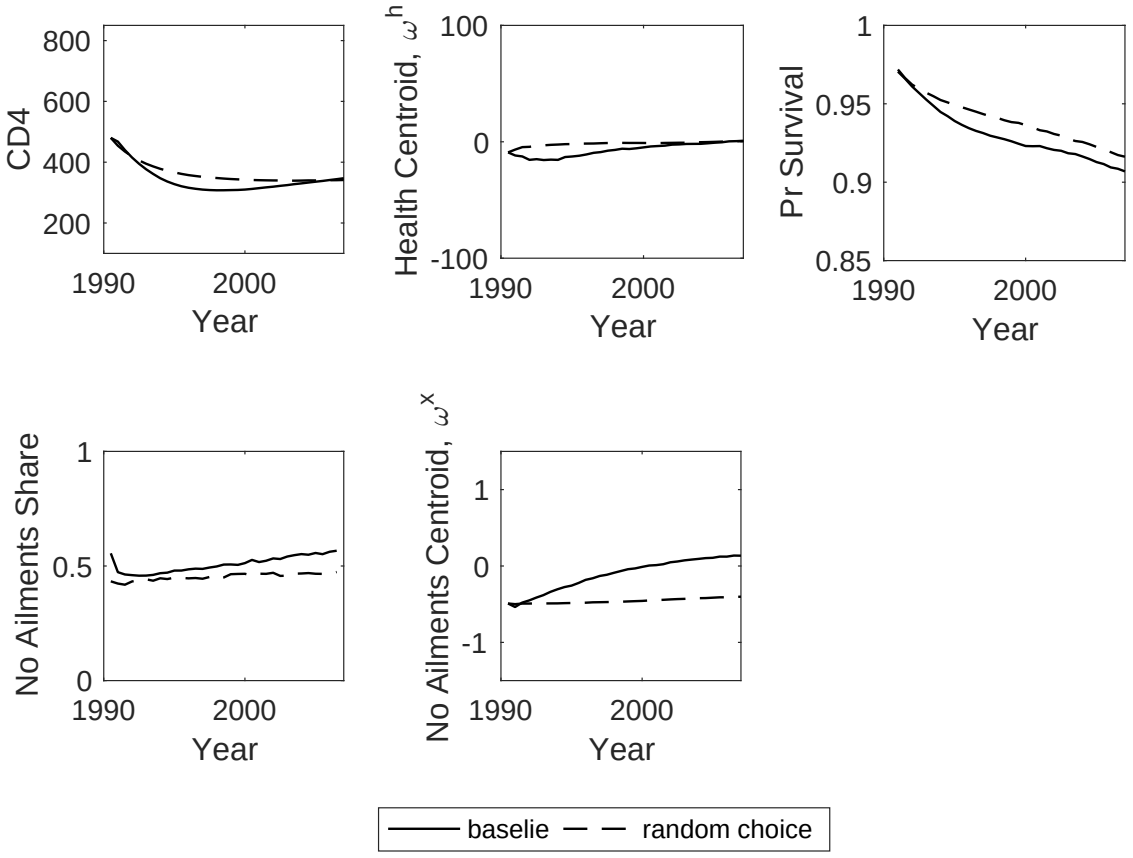

Figure 13: Alternative Regimes: Random Choice

Notes: Alternative choice regimes are: (i) baseline and (ii) random choice. The *baseline* regime is the estimated model where individuals make optimal dynamic choices. Individuals in the *random choice* regime select randomly using the aggregate shares on the baseline model to randomize. Averages computed from 500 simulations that are conditional on the state of the world at 1991.

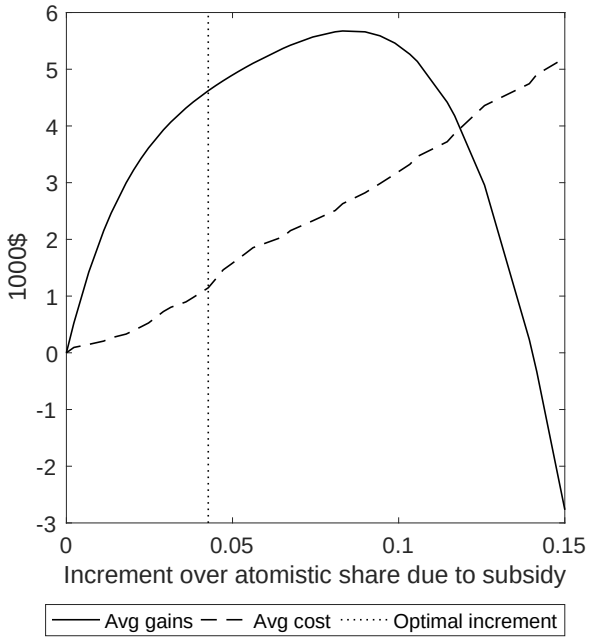

Figure 14: Optimal Assignment to Clinical Trials with a Flat Subsidy

Notes: The solid line represents the average welfare generated by a constrained planner in the second semester of 1996 that sends individuals to trials according to a rank based on how individuals like going to trials relative to their individually optimal choice. The dashed line indicates the average cost of using a flat subsidy to induce a given level of trial participation over the atomistic share. The dotted line indicates the planner's optimal share.

A Data Appendix

Beginning in 1984, the Multi-Center AIDS Cohort Study (MACS) started gathering information regarding natural and treated histories of HIV infection in homosexual and bisexual men. The study is conducted in Baltimore, Chicago, Pittsburgh and Los Angeles. At each semi-annual visit, data are collected on demographics, psychosocial characteristics, sexual behavior, and antiretroviral (AV henceforth) drugs consumption and trial participation. In addition, blood tests are administered to measure health status and serostatus (whether the individual is HIV+). Data collection started with 4,954 men enrolled. Two more enrollments have taken place: one in 1987-1991 (668 additional men) and another in 2001-2003 (1,350 additional men). We only use data from the first two enrollments. Since data is semi-annual each period t corresponds to 6 months.

Health (h_{it-1}): at every visit individuals undertake a physical examination that includes a blood sample which provides a measure of underlying health status: the individual's CD4 count. We denote as h_{it-1} the CD4 count at of the individual at the start of period t. According to the official U.S. government's website for HIV:[47]

> The CD4 count is [...] a snapshot of how well your immune system is functioning. CD4 cells (also known as CD4+ T cells) are white blood cells that fight infection. [...] These are the cells that the HIV virus kills. As HIV infection progresses, the number of these cells declines. When the CD4 count drops below 200 [cells per microliter] due to advanced HIV disease, a person is diagnosed with AIDS. A normal range for CD4 cells is about 500-1,500.

Labor supply (l_{it-1}): whether the individual worked full time (35 hours or more) during period t.

Income (m_{it}): starting at visit 14, individuals answer the question "*Which of the following categories describes your annual individual gross income before taxes?*" For visit 14, categories are brackets that increase every \$10,000 and are censored by the last category "\$70,000 or more." For visits 15 to 35 the brackets are censored at \$50,000 and for visits 36 to 41 the brackets are censored at \$60,000. We censor at \$50,000 to obtain a uniform question over time. Then we assign the middle point to individuals in the bracket. For the highest bracket we assign the upper limit (\$50,000). In our model gross income is divided by two since our periods are half-years. Gross income as well as out-of-pocket payments (below) are in constant dollars of 2000.

[47] See https://www.hiv.va.gov/patient/diagnosis/labs-CD4-count.asp

Out-of-pocket payments (o_{it}): starting at visit 14, individuals are asked a version of the following question "*Please, estimate the TOTAL out-of-pocket expenses that you or other personal sources (your lover, family or friends) paid for prescription medications since your last visit.*" This question is open so values are not categorized.

Ailments (x_{it}): starting at at visit 4, individuals are asked about physical symptoms. We focus on unusual bruises lasting at least two weeks, unintentional weight loss of at least 10 pounds, fatigue, diarrhea, fever, night sweats, and tender/enlarged glands. The last 5 ailments must be felt for at least 3 days during the period. Although individuals are asked explicitly about side effects starting at visit 13, we choose not to use this part of the data because it lacks consistency over time and more importantly, because individuals are most likely unable to correctly distinguish between side effects and symptoms. Thus, in our model x_{it} takes the value of 1 if an individual reports having any of the problems mentioned above.

Race (b_i) and *Age* (a_{it}): individuals are either white, black or hispanic, and their age at the beginning of period t, a_{it}, increases by half a year every period.

A.1 Products and Product Components

Starting at visit 6 individuals are asked about their medication. From visit 13 forward, as the number of treatments available increase, they answer separate survey modules for antiretroviral drugs (ARVs) and non antiretroviral drugs (NARVs). We focus on ARVs since these are the drugs used to treat HIV infection. Since our analysis entails estimating the health and ailments of people using different treatments, we focus on observations where individuals have reported a treatment along with their current and previous health (h_{it} and h_{it-1}), as well as their ailments (x_{it}).

Trial Products. Individuals are asked to name specifically which drugs they took as well as whether or not they took the drug as part of a research study. In the original data, some of the reported drugs are themselves coded as trials. We regard these instances as individuals participating in trials. If an individual consumes one of his drugs as part of a trial we regard the individual as consuming a trial product in that period.

Market Products. We define a market product as a combinations of components where no component is consumed in trial. For instance one product is AZT and another is AZT plus 3TC plus ddI. This definition results from noting that the sum of effects of consuming an individual drugs does not equal the effect of a treatment formed by the sum of the drugs because the interactions between components matter. Table S1 describes the individual components of market products and their usage. Some components in the original data are themselves fixed-dose combinations of other components (Table S2). In our sample, if

individual i is consuming the fixed-dose combination $(A+B)$ and individual i' is consuming components A and B, we assign consumers i and i' to the same treatment. Finally, we recode and add all uncoded components (96 observations) to the "Other ARVs" category.

Appendix Table S1: Product Components

Component	Chemical Formula	Observations
Isoprinosine	$C_{52}H_{78}N_{10}O_{17}$	87
Ribavirin	$C_8H_{12}N_4O_5$	62
Interferons (α and β)		210
Foscarnet	CH_3O_5P	92
AZT	$C_{10}H_{13}N_5O_4$	7436
ddC	$C_9H_{13}N_3O_3$	1123
AL-721 egg lecithin		147
Dextran-Sulfate	$H(C_6H_{10}O_5)_xOH$	65
Acyclovir	$C_8H_{11}N_5O_3$	2550
ddI	$C_{10}H_{12}N_4O_3$	3069
d4T	$C_{10}H_{12}N_2O_4$	3807
Nevirapine	$C_{15}H_{14}N_4O$	2210
Delavirdine	$C_{22}H_{28}N_6O_3S$	176
3TC	$C_8H_{11}N_3O_3S$	5250
Saquinavir	$C_{38}H_{50}N_6O_5$	1279
Ritonavir	$C_{37}H_{48}N_6O_5S_2$	3230
Indinavir	$C_{36}H_{47}N_5O_4$	2255
Nelfinavir	$C_{32}H_{45}N_3O_4S$	1278
Kaletra	$C_{37}H_{48}N_4O_5$	1883
Abacavir	$C_{14}H_{18}N_6O$	1549
Agenerase	$C_{25}H_{35}N_3O_6S$	372
Efavirenz	$C_{14}H_9CIF_3NO_2$	3362
Adefovir	$C_8H_{12}N_5O_4P$	44
Enfuvirtide (T-20)	$C_{204}H_{301}N_{51}O64$	160
Tenofovir	$C_9H_{14}N_5O_4P$	2488
Emtricitabine	$C_8H_{10}FN_3O_3S$	263
Atazanavir	$C_{38}H_{52}N_6O_7$	1583
Lexiva	$C_{25}H_{36}N_3O_9PS$	418
Etravirine	$C_{20}H_{15}BrN_6O$	155
Darunavir	$C_{27}H_{37}N_3O_7S$	315
Raltegravir	$C_{20}H_{21}FN_6O_5$	384
Ampligen	Double-stranded RNA compound	25
Peptide T	$C_{35}H_{55}N_9O_{16}$	30
DTC	$C_5H_{10}NS_2Na$	10
CD4		2
Other protease		31
Vistide (cidofovir)	$C_8H_{14}N_3O_6P$	2
Tipranavir (PNU-140690)	$C_5H_{10}NS_2Na$	30
Other ARVs		158

Notes: Source: Wikipedia (November, 2014). The *Observations* column indicates how many individuals in the sample used the combination component as part of a treatment.

Our definition of market products, as combinations of drug components, generates $1,835$ different treatments. We reduce the number of market products using the following algorithm:

Appendix Table S2: Combination Components

Name	Combination	Observations
Combivir	AZT + 3TC	2673
Trizivir	AZT + 3TC + Abacavir	778
Truvada	Emtricitabine + Tenofovir	1933
Epzicom	Abacavir + 3TC	724
Atripla	Efavirenz + Emtricitabine + Tenofovir	968

Notes: The *Observations* column indicates how many individuals in the sample used the drug as part of a treatment.

1. We select a set of "core market products" as those treatments that have more than 40 observations in the sample.[48] We acknowledge that our definition of core market products is biased against treatments appearing near the end of the time period studied. We address this issue by excluding the last 4 periods of data. Our core market products are listed in Table S3 which shows that there are 70 core market products overall and they have at most five components. Out of 20,142 subject-visit observations of consumers taking market products, 13,767 are covered by treatments classified as core market products.

2. We code the remaining 6,375 subject-visit observations of consumers taking a non-core market product as core market products. Each step sequentially assigns the remaining observations that were not assigned in previous steps.

   (a) Non-core market product A is assigned to core market product B if B is the core market product with the highest number of components that is contained by A. Of the remaining 6,375 subject-visit observations of non-core market products, this rule assigns 2,963 uniquely and leaves 3,412 with no unique assignment (1,647 that were assigned to multiple core market products plus 1,765 that were not assigned to any core market product).

   (b) If assigned to multiple core market products in Step (a):

      i. First, we use the past history of the individual. If at period t the individual is consuming non-core market product W that was assigned to both core market products A and B in Step (a), and he was observed consuming core market product A in period $t-1$, then his treatment at t is assigned uniquely as A. We repeat this procedure until no further gains are obtained. Out of the remaining 1,647 subject-visit observations assigned to multiple core market products, 428 are assigned uniquely in this step.

      ii. Second, we use the future history of the individual. If at period t the individual is consuming non-core market product W that was assigned to both

[48] We tried different minimum observations criteria and product classification did not change substantially.

core market products A and B in Step (a), and he was observed consuming core market product B in period $t+1$, then his treatment at t is assigned uniquely as B. We repeat this procedure until no further gains are obtained. Out of the remaining 1,219 subject-visit observations assigned to multiple core market products, 274 are assigned uniquely in this step.

iii. Third, we use the core market product with the highest share at t. If at period t the individual is consuming non-core market product W that was assigned to both core market products A and B in Step (a), and treatment A's market share at t is greater than B's, his treatment at t is assigned uniquely as A. This final step assigns uniquely the remaining 945 subject-visit observations assigned to multiple core market products.

(c) If not assigned to a core market product in Step (a): we regard all 1,765 subject-visit observations as "fringe treatments" since they do not contain any core market product. We aggregate all fringe treatments that appear at period t into one single "fringe mix," and assign to it all users consuming this product over time. We only consider fringe mixes that have at least 40 users. This reduces the number of observations by 345 (which represents 1.6% of the number of observations of treatment consumers). This aggregation leads to 16 fringe mixes that we pool with the set of core market products, which amounts to a total of 86 market products overall.

3. In the paper we specified that a treatment gets withdrawn from the market if it has zero share for 2 consecutive periods. However, in the data, a treatment may have zero share for more than 2 consecutive periods and then reappear again. 78 out of 86 core market products have unique spells without "reappearance." We regard the remaining treatments with multiple spells as measurement error and follow the next procedure to ensure that treatments have unique spells without reappearance. For every core market product B with reappearance:

(a) We identify all spells that treatment B has in the data. Notice that under this definition a single spell may contain some periods with zero share.

(b) From those spells we select the one that contains the period in which treatment B's share was the highest. We drop all observations of individuals consuming market product B in other spells.

Out of 19,797 (20,142-345) subject-visit observations of consumers taking market products, this smoothing procedure drops 42 observations leaving 19,755 subject-visit ob-

servations of consumers taking market products. As evidence of the relevance of the spells selected by this procedure the maximum share in the selected spell is on average about 24 times larger that the maximum share in other spells of the same market product.[49] Table S3 presents our market products as well as their entry and exit dates implied by this spell smoothing procedure.

B Model Appendix

B.1 k-means Clustering Algorithm

Our k-means clustering algorithm approximates the solution of the following objective function:[50]

$$\min_{1\{k\in g\}_{k\in\mathcal{P}_t}|\mathcal{G}_t} \quad \sum_{g=1}^{\mathcal{G}_t}\sum_{k\in\mathcal{P}_t} 1\{k\in g\}\left\|\tilde{\theta}_k-\tilde{\theta}_k^c\right\|^2$$
$$s.t.\sum_{g\in\mathcal{G}_t} 1\{k\in g\}=1 \text{ for all } k\in\mathcal{P}_t \tag{S1}$$

where the centroid of cluster k, $\tilde{\theta}_k^c$, is defined as

$$\tilde{\theta}_k^c=\frac{\sum_{k\in\mathcal{P}_t} 1\{k\in g\}\,\tilde{\theta}_k}{\sum_{k\in\mathcal{P}_t} 1\{k\in g\}} \tag{S2}$$

We implement the following version of the k-means algorithm. At every period t:

1. We select the products for which the *exit switching* rule has not been applied. In other words, we select products that are still available for people to switch into at period t. Denote this set of products available for clustering at t, $\mathcal{A}_t$.

2. In order to keep comparability we re-scale the characteristics of all products available for clustering at t by computing

$$\tilde{\theta}^r=\frac{\theta^r}{\max_{\delta\in\mathcal{A}_t}|\delta^r|}, \text{ for } r=h,x \tag{S3}$$

[49] We also tried (i) selecting the spell with the highest average share and (ii) selecting the spell with the highest sum of shares. All criteria result in very similar entry and exit dates so we stick to the maximum-share criteria.

[50] See Duda and Hart (1973) and Andrew W. Moore's *K-means and Hierarchical Clustering* tutorial at http://www.cs.cmu.edu/∼awm/tutorials.html.

Appendix Table S3: Market Products

Market Product	Haart	Entry	Exit	Market Product	Haart	Entry	Exit
AZT	0	1987 S1	-	ddI , d4T, Nevirapine	1	1997 S2	-
Interferons (α and/or β), AZT	0	1987 S2	1995 S2	ddI , 3TC, Nelfinavir	1	1997 S2	-
AL-721 egg lecithin	0	1987 S2	1991 S2	ddI , d4T, Efavirenz	1	1998 S2	2008 S1
AZT, Acyclovir	0	1989 S2	2000 S1	3TC, Abacavir, Efavirenz	1	1998 S2	-
Acyclovir	0	1989 S2	2000 S1	AZT, Nevirapine, 3TC, Abacavir	1	1999 S1	-
AZT, Acyclovir, ddI	0	1990 S1	1997 S1	AZT, 3TC, Abacavir, Efavirenz	1	1999 S1	-
Acyclovir, ddI	0	1990 S1	2000 S1	AZT, 3TC, Efavirenz	1	1999 S1	-
AZT, ddC	0	1990 S1	2001 S2	AZT, 3TC, Abacavir	0	1999 S1	-
AZT, ddI	0	1990 S1	2004 S2	d4T, 3TC, Efavirenz	1	1999 S1	2006 S1
ddI	0	1990 S1	-	Nevirapine, 3TC, Abacavir	1	1999 S2	-
AZT, ddC, Acyclovir, ddI	0	1991 S1	1997 S1	d4T, 3TC, Kaletra	1	2001 S1	2006 S1
AZT, ddC, Acyclovir	0	1991 S1	1999 S2	3TC, Kaletra, Abacavir	1	2001 S2	-
AZT, ddC, ddI	0	1991 S1	1995 S2	AZT, 3TC, Kaletra	1	2001 S2	-
ddC, Acyclovir	0	1991 S1	1997 S2	AZT, 3TC, Kaletra, Abacavir	1	2002 S1	-
ddC	0	1991 S1	1999 S1	3TC, Abacavir, Efavirenz, Tenofovir	1	2002 S1	-
d4T	0	1993 S1	-	AZT, 3TC, Abacavir, Tenofovir	1	2002 S1	-
AZT, Acyclovir, 3TC	0	1994 S2	2000 S1	AZT, 3TC, Kaletra, Tenofovir	1	2002 S1	-
AZT, 3TC	0	1995 S1	-	Nevirapine, 3TC, Tenofovir	1	2002 S1	2007 S1
Acyclovir, d4T, 3TC	0	1995 S2	2000 S1	3TC, Kaletra, Tenofovir	1	2002 S1	-
AZT, 3TC, Saquinavir	1	1996 S1	2005 S1	Kaletra, Efavirenz, Tenofovir	0	2002 S1	-
d4T, 3TC	0	1996 S1	-	3TC, Efavirenz, Tenofovir	1	2002 S1	-
AZT, 3TC, Saquinavir, Ritonavir	1	1996 S2	-	AZT, 3TC, Kaletra, Abacavir, Tenofovir	1	2002 S2	-
AZT, Acyclovir, 3TC, Indinavir	1	1996 S2	2000 S1	ddI , Kaletra, Tenofovir	1	2002 S2	-
Acyclovir, d4T, 3TC, Indinavir	1	1996 S2	2000 S1	ddI , Efavirenz, Tenofovir	1	2002 S2	-
AZT, 3TC, Ritonavir, Indinavir	1	1996 S2	2006 S2	Abacavir, Efavirenz, Tenofovir	1	2002 S2	-
d4T, 3TC, Ritonavir, Indinavir	1	1996 S2	2006 S2	Kaletra, Abacavir, Tenofovir	1	2002 S2	-
d4T, 3TC, Saquinavir, Ritonavir	1	1996 S2	2004 S2	3TC, Ritonavir, Abacavir, Atazanavir	1	2003 S2	-
ddI , d4T, Indinavir	1	1996 S2	2004 S2	Efavirenz, Tenofovir, Emtricitabine	1	2003 S2	-
d4T, 3TC, Indinavir	1	1996 S2	2008 S1	Ritonavir, Efavirenz, Tenofovir, Emtricitabine, Atazanavir	1	2004 S1	-
AZT, 3TC, Indinavir	1	1996 S2	-	3TC, Ritonavir, Abacavir, Tenofovir, Atazanavir	1	2004 S1	-
d4T, Nevirapine, 3TC	1	1997 S1	-	ddI , Ritonavir, Tenofovir, Atazanavir	1	2004 S1	-
AZT, Nevirapine, 3TC	1	1997 S1	-	Ritonavir, Tenofovir, Emtricitabine, Atazanavir	1	2004 S1	-
AZT, 3TC, Nelfinavir	1	1997 S1	-	Nevirapine, Tenofovir, Emtricitabine	1	2004 S1	-
ddI , d4T, Nelfinavir	1	1997 S1	2005 S2	Kaletra, Tenofovir, Emtricitabine	1	2004 S2	-
d4T, 3TC, Nelfinavir	1	1997 S2	-	Ritonavir, Tenofovir, Emtricitabine, Lexiva	1	2005 S1	-
				Fringe Mixes			
Isoprinosine, Ribavirin, Interferons (α and/or β)	0	1987 S1	1992 S1	Nevirapine, 3TC, Ritonavir, Kaletra, Tenofovir	0	2003 S1	-
Interferons (α and/or β), 3TC, Saquinavir, Indinavir, Efavirenz	0	1997 S1	2007 S1	3TC, Ritonavir, Kaletra, Abacavir, Tenofovir, Atazanavir	0	2004 S1	-
Nevirapine, 3TC, Saquinavir, Ritonavir, Indinavir	0	1997 S2	2006 S2	Ritonavir, Tenofovir, Emtricitabine, Atazanavir, Lexiva	1	2004 S2	-
Nevirapine, 3TC, Saquinavir, Ritonavir, Nelfinavir	0	1998 S1	2006 S2	Saquinavir, Ritonavir, Tenofovir, Emtricitabine, Atazanavir	1	2005 S1	-
Nevirapine, Saquinavir, Ritonavir, Abacavir, Efavirenz	1	1999 S1	2005 S2	3TC, Ritonavir, Abacavir, Tenofovir, Atazanavir, Lexiva	1	2005 S2	-
Nevirapine, Ritonavir, Nelfinavir, Abacavir, Efavirenz	0	1999 S2	-	Saquinavir, Ritonavir, Abacavir, Tenofovir, Emtricitabine	1	2007 S1	-
Nevirapine, Ritonavir, Kaletra, Abacavir, Efavirenz	0	2001 S2	2008 S2	3TC, Ritonavir, Tenofovir, Emtricitabine, Raltegravir	1	2008 S1	-
Nevirapine, 3TC, Nelfinavir, Abacavir, Tenofovir	1	2002 S2	-	Ritonavir, Tenofovir, Emtricitabine, Darunavir, Raltegravir	1	2008 S2	-

Notes: Entry and exit dates implied by the smoothing of spells in Step 3 of the algorithm used to reduce market products in Section A.1. S1 and S2 indicate the semester within a year. Many products had not exited by the end of the sample. The *Haart* column indicates whether a market product is a member of the Highly Active Antiretroviral Treatment class. For *Fringe Mixes* we only include the 5 or 6 most used products in the mix.

Thus, by construction $\tilde{\theta}^r \in [-1, 1]$.

3. We select the first k centroids using the scaled characteristics vectors $\tilde{\theta}$ of k randomly selected products from $\mathcal{A}_t$.

4. We allocate all remaining products in $\mathcal{A}_t$ to clusters sequentially. At each step the product selected for allocation is the one whose scaled characteristics $\tilde{\theta}$ are closest to one of the existing clusters. This point is then allocated to the closest cluster and the centroid of the cluster is updated. This process is repeated until all points are allocated to a cluster.

5. We undertake a reallocation step in which, taken the centroids as given, all points are allocated to their closest centroid.

6. We calculate the value of (S1) for the current allocation.

7. We repeat 200 times Steps 3 to 6 using the scaled characteristics $\tilde{\theta}$ of different groups of k randomly selected products in $\mathcal{A}_t$ as initial centroids. The allocation with the lowest value of (S1) is chosen. In estimation, whenever we simulate clusters we only repeat the process 50 times.

B.2 Timing

The aggregate state in period t is denoted Ξ_t and consists of current and previous product characteristics $\mathbb{P}_t$, previous market shares $\mathbb{S}_t$, and the distribution of current-period consumer characteristics $\mathcal{F}_t$. Together, these factors determine a summary of the state of technology ω_t, entry (the number of new products and their characteristics), and exit of products that are withdrawn from the market prior to the start of next period.

The consumer observes his individual state $\mathcal{Z}_{it}$, along with his choice set $\mathcal{C}_t$, upon entering period t. His state consists of individual-level components $z_{it}^{\mathcal{I}}$ (e.g., health) along with aggregate market components $z_{it}^{\mathcal{M}}$. When choosing an alternative, he takes account of how each choice affects current outcomes (e.g., side effects, income) and future states (e.g., health, labor participation). He also forecasts the characteristics of future alternatives, which may affect the relative payoffs to his current choice. The consumer's choice of treatment maximizes his expected discounted lifetime utility. Once a consumer makes a choice, outcome variables are realized and he receives his flow utility. Thereafter, their state variables update and the consumer enters the next period. The timing of the problem is illustrated in Figure S1.

Appendix Figure S1: Timing

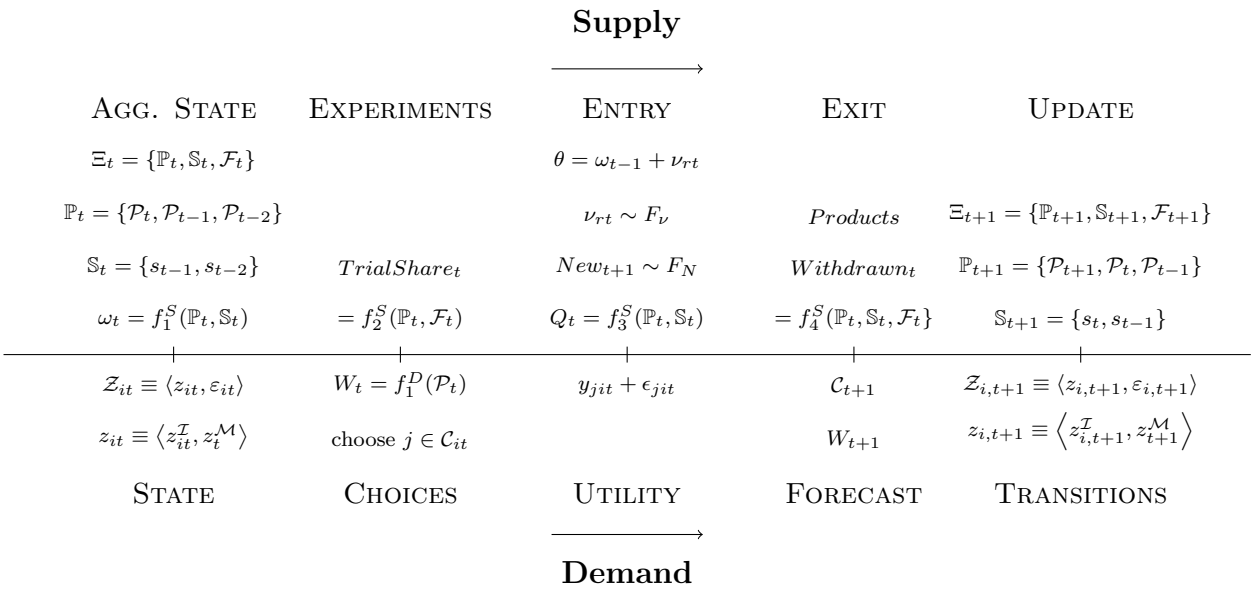

C Estimation Appendix

C.1 Product Characteristics

We estimate market product characteristics using data on individual treatment usage and subsequent health and ailments. The estimation equations mimic equations (19) and (20), which individuals use to form expectations over their health and ailments conditional on their choice. The key difference is that here our aim is to obtain characteristics of each market product. Let δ_{rit} be an indicator that market product r was used by individual i at time t. The characteristics of market product r are denoted

$$\theta_r = \left\{\theta_r^x, \theta_r^h\right\} \in \mathbb{R}^2 \tag{S4}$$

The components of θ_r are estimated as the coefficients of δ_{rit} in the health and no-ailments regressions

$$\Pr\left[x_{it} = 1|\cdot\right] = \frac{\exp\left(\sum_{m=0}^{5} \gamma_m^x h_{it-1}^m + \sum_r \theta_r^x \delta_{rit}\right)}{1 + \exp\left(\cdot\right)} \tag{S5}$$

$$h_{it} = \sum_{m=0}^{5} \gamma_m^h h_{it-1}^m + \sum_r \theta_r^h \delta_{rit} + \epsilon_{it} \tag{S6}$$

Estimation of equations (S5) and (S6) also provides the vectors of parameters γ^x and γ^h to be used in equations (19) and (20).

C.2 Proof of Proposition 1

The proof of Proposition 1 boils down to the representation of the conditional value function v_{jit} in terms of future utility flows induced by all available choices, weighted by the future probabilities of those choices and corrected by the fact that the choice may not be optimal. Notably, the weighted average of corrected flow payoffs of a given period is discounted by the probability of survival up to that period conditional on today's state and choice, as well as conditional on optimal behavior. To decrease notation baggage, drop the individual subindex i. Let d^o denote optimal behavior, d^o_{kt} be an indicator for choice k being optimal at t, and let $S_{t+r}(z_{t+r})$ and $D_{t+r}(z_{t+r})$ be as in equation (21). The conditional value function is given by

$$
\begin{aligned}
v_{jt}(z_t) &= \mathbb{E}_y[y_{jt}|z_t] + \beta\mathbb{E}\left[V(z_{t+1},\varepsilon_{t+1})\,|z_t,j\right] \\
&= \mathbb{E}_y[y_{jt}|z_t] + \beta\mathbb{E}\left[S_{t+1}(z_{t+1})\,\mathbb{E}_\varepsilon\left[\sum_{k\in C_{t+1}} d^o_{kt+1}(z_{t+1})\left[y_{kt+1}(z_{t+1})+\psi_{kt+1}(z_{t+1})\right]\right]\middle|\, z_t,j\right] \\
&\quad +\beta^2\mathbb{E}\left[S_{t+1}(z_{t+1})\,V(z_{t+2},\varepsilon_{t+2})\,|z_t,j,d^o\right] \\
&= \mathbb{E}_y[y_{jt}|z_t] + \beta\mathbb{E}\left[S_{t+1}(z_{t+1})\sum_{k\in C_{t+1}} p_{kt+1}(z_{t+1})\left[y_{kt+1}(z_{t+1})+\psi_{kt+1}(z_{t+1})\right]\middle|\, z_t,j\right] \\
&\quad +\beta^2\mathbb{E}\left[S_{t+1}(z_{t+1})\,V(z_{t+2},\varepsilon_{t+2})\,|z_t,j,d^o\right] \\
&= \mathbb{E}_y[y_{jt}|z_t] + \beta\mathbb{E}\left[S_{t+1}(z_{t+1})\sum_{k\in C_{t+1}} p_{kt+1}(z_{t+1})\left[y_{kt+1}(z_{t+1})+\psi_{kt+1}(z_{t+1})\right]\middle|\, z_t,j\right] \\
&\quad +\beta^2\mathbb{E}\left[S_{t+1}(z_{t+1})\,S_{t+2}(z_{t+2})\sum_{k\in C_{t+1}} p_{kt+2}(z_{t+2})\left[y_{kt+2}(z_{t+2})+\psi_{kt+2}(z_{t+2})\right]\middle|\, z_t,j,d^o\right] \\
&\quad +\beta^3\mathbb{E}\left[S_{t+1}(z_{t+1})\,S_{t+2}(z_{t+2})\,V(z_{t+3},\varepsilon_{t+3})\,|z_t,j,d^o\right] \\
&= \mathbb{E}_y[y_{jt}|z_t] + \sum_{s=1}^{T^*}\beta^s\mathbb{E}\left[\left(\prod_{r=1}^{s} D_{t+r}(z_{t+r})\right)\sum_{k\in C_{t+1}} p_{kt+s}(z_{t+s})\left[y_{kt+s}(z_{t+s})+\psi_{kt+s}(z_{t+s})\right]\middle|\, z_t,j,d^o\right] \\
&\quad +\beta^{T^*+1}\mathbb{E}\left[\left(\prod_{r=1}^{T^*} D_{t+r}(z_{t+r})\right)V(z_{t+T^*+1},\varepsilon_{t+T^*+1})\middle|\, z_t,j,d^o\right]
\end{aligned}
\tag{S7}
$$

That

$$
\psi_{kit}(z_{it}) = \gamma - \ln\left(p_{kit}(z_{it})\right) \tag{S8}
$$

follows from the joint distribution of the taste shifter ε_{it}, which is Extreme Value Type-I. *Q.E.D.*

C.3 GMM Estimation

In order to form the sample analog of the moment condition in (26) we use the results from Proposition 1 to obtain a simulated version of the conditional value function truncated at T^* for each individual i and choice j at every period t. The truncated conditional value function is

$$\mathbb{E}_y[y_{jit}|z_{it}] + \sum_{s=1}^{T^*} \beta^s \mathbb{E}\left[\left(\prod_{r=1}^{s} D_{it+r}\left(z_{it+r}\right)\right) \sum_{k \in C_{it+1}} p_{kit+s}\left(z_{it+s}\right)\left[y_{kit+s}\left(z_{it+s}\right) + \psi_{kit+s}\left(z_{it+s}\right)\right] \middle| z_{it}, j, d_i^o\right] \tag{S9}$$

Let NS denote the number of simulated technology paths for each observation (subject-visit) and let the superscript ns be a simulation index. For individual i and decision j at period t we write the simulated counterpart of equation (S9) as

$$\frac{1}{NS}\sum_{ns}\left\{\mathbb{E}_y[y_{jit}^{ns}|z_{it}] + \sum_{s=1}^{T^*} \beta^s \left[\left(\prod_{r=1}^{s} D_{it+r}\left(z_{it+r}^{ns}\right)\right) \sum_{k \in C_{it+1}^{ns}} d_{kit+s}^{ns}\left(z_{it+s}^{ns}\right)\left[y_{kit+s}^{ns}\left(z_{it+s}^{ns}\right) + \psi_{kit+s}^{ns}\left(z_{it+s}^{ns}\right)\right]\right]\right\} \tag{S10}$$

The forward simulation depends on the current individual state z_{it}, the current aggregate state Ξ_{it}, and the current choice j. For a given vector of parameters of the utility function, simulation must be undertaken NS times for each individual i available at period t for all J choices, which amounts to $NS \times T \times N \times J$ simulations. Further, notice that within each of those individual simulations we must simulate optimal paths for all N individuals in order to obtain the aggregate behavior and technological paths. In other words, we must simulate $NS \times T \times N \times J \times N$ individual paths. Given our data this amounts to about $NS \times 33 \times 1669 \times 6 \times 1669 = NS \times 551,541,078$ individual paths of length T^*. Because this is too computationally taxing, we instead implement the following procedure. We first simulate one path of aggregate behavior unique to every observation (subject-visit) in the data, which allows us to construct as many paths of technological innovation as there are observations. Notice that this entails forward simulating the choices of all individuals in the sample at period t for every individual observed at t. Then, because individuals are atomistic, for each observation i at period t and choice j we generate sequences of future choices and payoffs taking as given a subset of $NS = 20$ artificial technological paths chosen at random from the already generated subset of technological paths that start at date t.[51]

Let $\bar{v}_{kit}\left(z_{it}\right)$ denote the truncated simulated conditional value function in equation (S10), let o denote a base choice, and let δ_{it} be an indicator of whether individual i is in the data

[51] Notice that we could rely on Hotz et al. (1994) and set $NS = 1$ and still obtain consistency. However, we choose $NS = 20$ after trying different values for robustness.

at period t. The sample analog of the moment condition in (26) is

$$\frac{1}{\sum_i \sum_t \delta_{it}} \sum_{i=1}^{N} \sum_{t=1}^{T} \delta_{it} w\left(z_{it}\right) \otimes \begin{bmatrix} \vdots \\ \ln\left(\frac{p_{oit}(z_{it})}{p_{1it}(z_{it})}\right) + \bar{v}_{1it}(z_{it}) - v_{oit}(z_{it}) \\ \vdots \\ \ln\left(\frac{p_{oit}(z_{it})}{p_{J-1it}(z_{it})}\right) + \bar{v}_{J-1it}(z_{it}) - v_{oit}(z_{it}) \\ \vdots \end{bmatrix} = 0 \tag{S11}$$

C.3.1 Simulation of Aggregate Paths

In order to obtain $\bar{v}_{kit}\left(z_{it}\right)$ we first create the set of simulated technological paths following the steps bellow for each individual i at period t taking as given their observed choice j, their current state z_{it}, and the current aggregate state Ξ_{it}:

1. Let $s = 1$.

2. *Number of new products.* Simulate a number of new products at $t+s$, New_{t+s}^{ns}, using the stochastic process for entry in (5).

3. *Characteristics of new products.* If $New_{t+s}^{ns} > 0$, for each simulated new product draw simulated product characteristics using equations (2) and (3). As a by-product of steps 1 and 2 obtain Q_{t+s}^{ns} using equation (6).

4. *Exit.* For all incumbent products, apply both exit rules described in Section 3 (*overall exit* and *exit for switchers*). Old products minus exits plus simulated new products yields the simulated set of products in $t+s$, $\mathcal{P}_{t+s}^{ns}$.

5. *Clusters.* From the simulated set of products $\mathcal{P}_{t+s}^{ns}$ that have not exited for switchers, form clusters $\mathcal{G}_{t+s}^{ns}$ following the clustering algorithm explained in Appendix Section B. Compute moments of the simulated clusters, denoted W_{t+s}^{ns}.

6. *Centroid.* If $s = 1$, $\mathcal{P}_{t+s-1}^{ns} \equiv \mathcal{P}_t$. Compute the simulated centroid ω_{t+s}^{ns} using equation (1). Steps 2 through 6 provide the aggregate part of the simulated state, $z_{t+s}^{\mathcal{M},ns}$. Denote the future choice set induced by the simulated evolution of products as $\mathcal{C}_{t+s}^{ns}$.

7. *Future individual state.* For all individuals i' at t: **(i)** If $s = 1$, define $h_{i't+s-1}^{ns} \equiv h_{i't+s-1}$ and $d_{i't+s-1}^{ns} \equiv d_{i't+s-1}$. If $1 < s < T^*$, simulate health at the beginning of $t+s$, $h_{i't+s-1}^{ns}$, using equation (20). This entails drawing a health shock $\epsilon_{i't+s-1}^{h,ns}$ and using the previous choice $d_{i't+s-1}^{ns}$ and when necessary the realization of the within cluster

treatment assigned at $t+s-1$. If $d^{ns}_{i't+s-1}$ is the trial alternative, draw trial-product characteristics for computing equation (20) using equations (2) and (3). **(ii)** Draw a simulated labor state $l^{ns}_{i't+s}$ using equation (17). **(iii)** Compute deterministic transitions (e.g. age). This step yields the individual-specific part of the simulated state $z^{\mathcal{I},ns}_{i't+s}$.

8. *CCPs and simulated choice.* For all individuals i' at t: using $z^{\mathcal{M},ns}_{t+s}$, $z^{\mathcal{I},ns}_{i't+s}$, and equations (S26), (S27) and (S28) compute simulated ccps $p^{ns}_{ki't+s}\left(z^{ns}_{i't+s}\right)$ for every alternative $k \in \mathcal{C}^{ns}_{t+s}$ and draw a decision $d^{ns}_{i't+s}\left(z^{ns}_{i't+s}\right)$.

9. *Cycle back.* If $s = T^*$ end the loop. Otherwise, let $s = s+1$ and go back to Step 2.

As a result from this algorithm we obtain as many simulated technological paths as there are observations in the data. Denote the set of simulated paths starting at t as $\mathcal{T}_t$.

C.3.2 Simulation of Individual Paths

Since individuals are atomistic and we want to minimize simulation error, for each individual i at period t and choice j we generate NS sequences of future choices and payoffs taking as given a subset of $NS = 20$ artificial technological paths chosen at random from $\mathcal{T}_t$. Hence, for each individual i at period t for choice j and artificial technological path ns we take the following steps:

1. Let $s = 1$.

2. *Future individual state.* Same as above but only for individual i. When j is not equal to the observed choice for individual i at period t a state must be simulated when $s = 1$ as well.

3. *Survival probability.* Compute the simulated probability of surviving up to $t+s-1$ as[52]

$$\prod_{r=1}^{s-1} D_{it+r}\left(z^{ns}_{it+r}\right) \tag{S12}$$

4. *CCPs and simulated choice.* Same as above but only for individual i.

[52] We do not simulate dead and instead weight the individual future payoffs by his simulated probability of survival.

5. *Static payoff.* **(i)** Simulate expected income using equation (15).[53] **(ii)** Simulate expected out-of-pocket payments using equation (16).[54] **(iii)** Simulate the expected probability of no-ailments using equation (19) and the relevant distribution of product characteristics implied by the simulated choice d^{ns}_{it+s}. For instance, whenever the choice is a cluster, use the within cluster weights. **(iv)** Compute expected flow payoffs $y^{ns}_{it+s}\left(z^{ns}_{it+s}, d^{ns}_{it+s}\right)$ using equation (13). **(v)** Compute the correction term $\psi_{it+s}\left(z^{ns}_{it+s}, d^{ns}_{it+s}\right)$ using equation (25).

6. *Cycle back.* If $s = T^*$ end the loop. Otherwise, let $s = s + 1$ and go back to Step 2.

C.3.3 Initial Counterfactual State and Payoff

When simulating the path following choice $j' \neq j$, where j is the observed choice at t, we also need to simulate health at the beginning of period $t + s$ for $s = 1$ and current payoffs. We back out the realized health residuals using equation (S6) and use equation (20) to simulate health h^{ns}_{it}. If individual i was in a trial in period t we do not observe the characteristics of the trial product ex post, so we draw a health shock as well as trial product characteristics to simulate health.

In order to obtain current-period simulated expected payoffs $E_y[y^{ns}_{j'it}|z_{it}]$ for counterfactual choice j' we need expected income, out-of-pocket payments, and ailments. Whenever j' corresponds to a cluster alternative we use the within cluster distribution of characteristics to compute expected ailments. For out-of-pocket payments we need the realized error term of the out-of-pocket payment equation (16) at t given by

$$\hat{\epsilon}^{o}_{it} = o^{*}_{jit} - X^{o}_{jit}\theta^{o} \tag{S13}$$

However, we only observe o^{*}_{jit} if $o^{*}_{jit} > 0$. Hence, if $o^{*}_{jit} \leq 0$, we draw a simulated error $\epsilon^{o,ns}_{it}$ from a truncated normal conditional on $\epsilon^{o,ns}_{it} \leq -X^{o}_{jit}\theta^{o}$.

[53] Even though individuals know their idiosyncratic income shocks ϵ^{m}_{it} we do not need to simulate these shocks as they are iid, have mean zero, and enter linearly in the flow utility, which results in them averaging out to zero in the moment condition.

[54] Defining $X^{o,ns}_{it+s}\left(d^{ns}_{it+s}\right)$ as in equation (16):

$$\mathbb{E}\left[o_{it+s}\left(d^{ns}_{it+s}\right)\middle|\, d^{ns}_{it+s}\right] = \Phi\left(X^{o,ns}_{it+s}\left(d^{ns}_{it+s}\right)\theta^{o}/\sigma^{o}\right) X^{o,ns}_{it+s}\left(d^{ns}_{it+s}\right)\theta^{o} + \sigma^{o}\phi\left(X^{o,ns}_{it+s}\left(d^{ns}_{it+s}\right)\theta^{o}/\sigma^{o}\right)$$

C.3.4 Estimator

We use a GMM estimator to obtain the parameters of the utility function. Define B as the $K-$dimensional vector of parameters of the utility function. Following Hotz et al. (1994) we estimate B as the vector that maximizes the following objective function

$$\left((NT)^{-1}\sum_{i=1}^{N}\sum_{t=1}^{T}\delta_{it}w\left(z_{it}\right)\otimes A_{it}(z_{it},B)\right)' W_n \left((NT)^{-1}\sum_{i=1}^{N}\sum_{t=1}^{T}\delta_{it}w\left(z_{it}\right)\otimes A_{it}(z_{it},B)\right) \tag{S14}$$

$$A_{it}(z_{it},B)\equiv\begin{bmatrix} \vdots \\ \ln\left(\frac{p_{oit}(z_{it})}{p_{1it}(z_{it})}\right)+\bar{v}_{1it}(z_{it})-v_{oit}(z_{it}) \\ \vdots \\ \ln\left(\frac{p_{oit}(z_{it})}{p_{J-1it}(z_{it})}\right)+\bar{v}_{J-1it}(z_{it})-v_{oit}(z_{it}) \\ \vdots \end{bmatrix} \tag{S15}$$

where W_n is a square weighting matrix. Using the linear structure of the utility function in equation (13) we collect and factor terms in order to write the jth component of the vector $A_{it}(z_{it},B)$ as the linear form

$$\tilde{y}_{jit}-\tilde{x}'_{jit}B \tag{S16}$$

Define Y as a vector with $(J-1)NT$ rows that stacks all $\tilde{y}_{jit}$, and X as a $(J-1)NT\times K$ matrix that stacks all $\tilde{x}_{jit}$. Define Z as the $NT\times R$ matrix whose columns contain the R instruments orthogonal to the difference in alternative representations of the conditional value functions.[55] Thus

$$Y=\begin{bmatrix} \tilde{y}_{1,1,1} \\ \tilde{y}_{1,1,2} \\ \vdots \\ \tilde{y}_{1,N,T-1} \\ \tilde{y}_{1,N,T} \\ \vdots \\ \tilde{y}_{J-1,1,1} \\ \tilde{y}_{J-1,1,2} \\ \vdots \\ \tilde{y}_{J-1,N,T-1} \\ \tilde{y}_{J-1,N,T} \end{bmatrix},\quad X=\begin{bmatrix} \tilde{x}_{1,1,1,1} & \cdots & \tilde{x}_{1,1,1,K} \\ \tilde{x}_{1,1,2,1} & \cdots & \tilde{x}_{1,1,2,K} \\ \vdots & & \vdots \\ \tilde{x}_{1,N,T-1,1} & \cdots & \tilde{x}_{1,N,T-1,K} \\ \tilde{x}_{1,N,T,1} & \cdots & \tilde{x}_{1,N,T,K} \\ \vdots & & \vdots \\ \tilde{x}_{J-1,1,1,1} & \cdots & \tilde{x}_{J-1,1,1,K} \\ \tilde{x}_{J-1,1,2,1} & \cdots & \tilde{x}_{J-1,1,2,K} \\ \vdots & & \vdots \\ \tilde{x}_{J-1,N,T-1,1} & \cdots & \tilde{x}_{J-1,N,T-1,K} \\ \tilde{x}_{J-1,N,T,1} & \cdots & \tilde{x}_{J-1,N,T,K} \end{bmatrix},\quad Z=\begin{bmatrix} w\left(z_{11}\right)_1 & \cdots & w\left(z_{11}\right)_R \\ w\left(z_{12}\right)_1 & \cdots & w\left(z_{12}\right)_R \\ \vdots & & \vdots \\ w\left(z_{NT}\right)_1 & \cdots & w\left(z_{NT}\right)_R \end{bmatrix} \tag{S17}$$

[55] Hence W_n is a $(J-1)R$-dimensional square matrix.

Finally, let $\mathbf{I}_{[J-1]}$ be a $(J-1)$-dimensional identity matrix and define

$$\tilde{Z} = \mathbf{I}_{[J-1]} \otimes Z \tag{S18}$$

Then we can write the objective function in (S14) as

$$\left((NT)^{-1}\tilde{Z}'(Y - XB)\right)' W_n \left((NT)^{-1}\tilde{Z}'(Y - XB)\right) \tag{S19}$$

Equation (S19) is a linear arrangement so we can obtain a close form solution for $\hat{B}$ as the optimal GMM estimator. It entails a first stage estimator given by

$$\hat{B}^{1S} = \left(X'\tilde{Z}\tilde{Z}'X\right)^{-1}\left(X'\tilde{Z}\tilde{Z}'Y\right) \tag{S20}$$

and a second stage estimator given by

$$\hat{B}^{2S} = \left(X'\tilde{Z}\hat{S}^{-1}\tilde{Z}'X\right)^{-1}\left(X'\tilde{Z}\hat{S}^{-1}\tilde{Z}'Y\right) \tag{S21}$$

where

$$\hat{S} = \frac{1}{N^*}\tilde{Z}'D\tilde{Z} \tag{S22}$$

and D is the $N(J-1)$ square diagonal matrix with diagonal elements $\hat{u}^2_{jit} = \left(y_{jit} - x'_{jit}\hat{B}^{1S}\right)^2$. The variance-covariance matrix of the second stage estimator is

$$\hat{V}^{2S} = N^*\left(X'\tilde{Z}\hat{S}^{-1}\tilde{Z}'X\right)^{-1} \tag{S23}$$

and

$$N^* = \sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{j=1}^{J-1} 1\{\text{decision } j \text{ available for } i \text{ at } t\} \tag{S24}$$

which accounts for the fact that some individuals cannot stay in their lagged treatments at some periods (for instance, if their lagged choice was no treatment or a trial product).

We use as instruments lagged health h_{it-1}, lagged labor state l_{it-1}, income fixed effect η_i, race and education indicators, age a_{it-1}, the centroid ω_t and the lagged share of trial participation, as well as interactions between these variables.

C.4 Standard Errors

The uncorrected standard errors for our utility parameters yield from the variance-covariance matrix in equation S23. In order to obtain corrected standard errors we undertake subsam-

pling taking as given the following objects obtained from the full sample: the definition of products (i.e. what their components are, for instance, AZT or AZT + ddI), their corresponding entry and exit dates, and the exit thresholds $\tilde{\sigma}_1$ and $\tilde{\sigma}_2$ specified in Section 3. We draw $R = 100$ subsamples containing a proportion $\tilde{p} = 0.9$ of the individuals in the sample drawn without replacement, and estimate all parameters in the model using each subsample. This includes estimating product characteristics, parameters governing transition and outcome processes, and simulating forward paths of technology to obtain utility parameters. For any parameter γ with estimated value $\hat{\gamma}_r$ from the rth subsample, the subsampling standard errors are obtained as

$$se(\hat{\gamma}) \approx se(\hat{\gamma}_r) \cdot \sqrt{\tilde{p}} \tag{S25}$$

where $se(\hat{\gamma}_r)$ is estimated as the standard deviation of the R quantities $\hat{\gamma}_r$.

C.5 Estimated CCPs

The probability that an individual chooses one of the alternatives depends on the individual and aggregate elements of his state. Individuals decide between one of $\mathcal{G}$ clusters, yesterday's product (if any), a trial product, and no product. Let W_{jit} be the characteristics describing alternative j for individual i at period t: mean health, mean ailments, and the variance matrix. Let $W_{jit}W_{jit}$ denote a vector of interactions between the elements of W_{jit}. Let $\tilde{x}_{it}$ and $\tilde{z}_{it}$ be subsets of the individual-specific components of the state.[56] Let $\omega_t W_{jit}$ denote a vector of interactions between the centroid and the elements of W_{jit}. Similarly, let $W_{jit}\tilde{z}_{it}$ be a vector of interactions between the components of W_{jit} and individual-specific state components and let $\omega_t W_{jit}\tilde{z}_{it}$ be defined in a similar fashion. Finally, let $\tilde{\mathcal{F}}_t$ denote a set of non parametric moments describing the joint distribution of aggregate characteristics, $\mathcal{F}_t$.[57] For each of the alternatives, the CCPs are expressed as follows:

Cluster ccps $(j = 1, \ldots, \mathcal{G})$

$$p_{jit} = \frac{\exp\left(\gamma_0\tilde{x}_{it} + \beta_0 W_{jit} + \beta_1 W_{jit}W_{jit} + \beta_2\omega_t W_{jit} + \beta_3 W_{jit}\tilde{z}_{it} + \beta_4\omega_t W_{jit}\tilde{z}_{it} + \beta_5 W_{jit}\tilde{\mathcal{F}}_t\right)}{1 + \sum_{k=1}^{\mathcal{G}+2}\exp(\cdot)} \tag{S26}$$

γ_0 is constant across clusters and over time. For a given cluster j and period t, W_{jit} is in fact constant across individuals so $W_{jit} = W_{jt}$.

[56] $\tilde{z}_{it}$ includes h_{it-1}, a_{it-1}, b_i, l_{it} while $\tilde{x}_{it}$ includes a constant, a_{it-1}, b_i.

[57] We specify these moments as shares of people with different sets of characteristics.

Trial ccps ($j = \mathcal{G} + 1$)

$$p_{jit} = \frac{\exp\left(\gamma_j \tilde{x}_{it} + \beta_0 W_{jit} + \beta_1 W_{jit} W_{jit} + \beta_3 W_{jit} \tilde{z}_{it} + \beta_5 W_{jit} \tilde{\mathcal{F}}_t\right)}{1 + \sum_{k=1}^{\mathcal{G}+2} \exp(\cdot)} \tag{S27}$$

For the trial alternative, W_{jit} is constant across individuals so $W_{\mathcal{G}+1it} = W_{\mathcal{G}+1t}$. Since some of the components of W_{jt} are linear functions of $\omega_{t-1} + \mu_\nu$ (see equation (2)) we avoid collinearity by not including terms $\omega_t W_{jt}$ and $\omega_t W_{jt} \tilde{z}_{it}$ in the trials ccps.

Staying ccps ($j = \mathcal{G} + 2$)

$$p_{jit} = \frac{\exp\left(\gamma_j \tilde{x}_{it} + \beta_0 W_{jit} + \beta_1 W_{jit} W_{jit} + \beta_2 \omega_t W_{jit} + \beta_3 W_{jit} \tilde{z}_{it} + \beta_4 \omega_t W_{jit} \tilde{z}_{it} + \beta_5 W_{jit} \tilde{\mathcal{F}}_t\right)}{1 + \sum_{k=1}^{\mathcal{G}+2} \exp(\cdot)} \tag{S28}$$

When individuals choose to stick to their previous product $W_{\mathcal{G}+2it}$ becomes heterogeneous as individuals may have consumed different products last period.

No product ccps ($j = 0$)

$$p_{jit} = 1 - \sum_{k=1}^{\mathcal{G}+2} p_{kit} \tag{S29}$$

Although the characteristics of the choice set are non stationary, by interacting our time-varying regressors $\tilde{z}_{it}$ with the characteristics of the choice for individual i, W_{jit}, we are able to control for the state of the world inside the ccps. By following this approach avoid estimating period-specific logits for the ccps. Notably, this procedure gives us ccps for any simulated world as long as our observed worlds cover the space of possible worlds reasonably well. Additionally, we include in the ccps parameters that are invariant to the state of the technology, denoted γ, which capture stationary taste differences between staying in a choice (when possible), trying a new market product, going to a trial, or not consuming anything. Also, since conditional on cluster characteristics all clusters are equivalent to "trying a new market product" we impose $\gamma_j = \gamma_{j'} = \gamma_0$ for any $j, j' = 1, \ldots, \mathcal{G}$. Figures S2, S4, and **??** display the mean predicted conditional choice probability using equations (S26) to (S29) over time against the correspondent share of the population who chose the alternative. Our ccps map the choices in the data fairly well.[58]

[58] We further explore the fit of our ccp estimates comparing the relatives shares that clusters received in reality against the predictions from our estimated ccps. We ranking the three clusters at every period by the share they received and compare this ranking against the ranking obtained from our estimated ccps. A cross tabulation of these rankings suggests that the predicted ranks match the real ranks in about 80% of the periods.

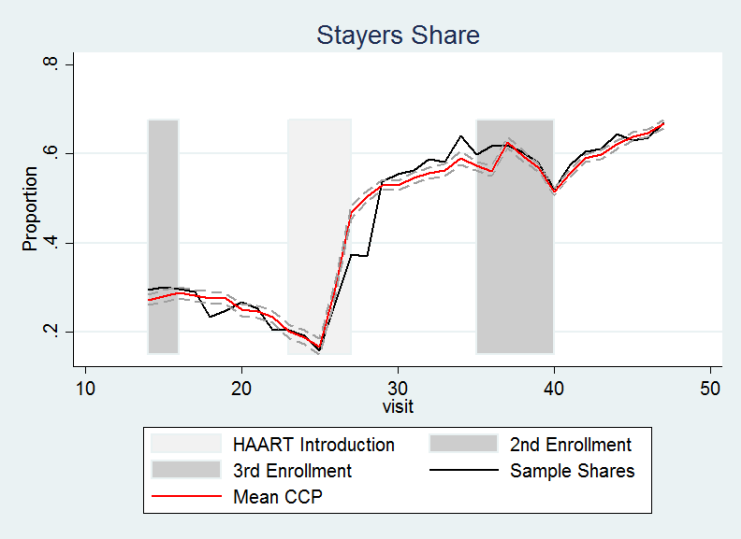

Appendix Figure S2: Average CCPs for Staying

Notes: The figure shows the average estimated conditional choice probability against the share of people choosing the alternative. Dashed lines represent 95% confidence intervals around the predicted CCPs. Three periods of special relevance are highlighted in the Figure: two periods during which enrollment into the sample was undertaken and the period in which products belonging to the HAART class were introduced.

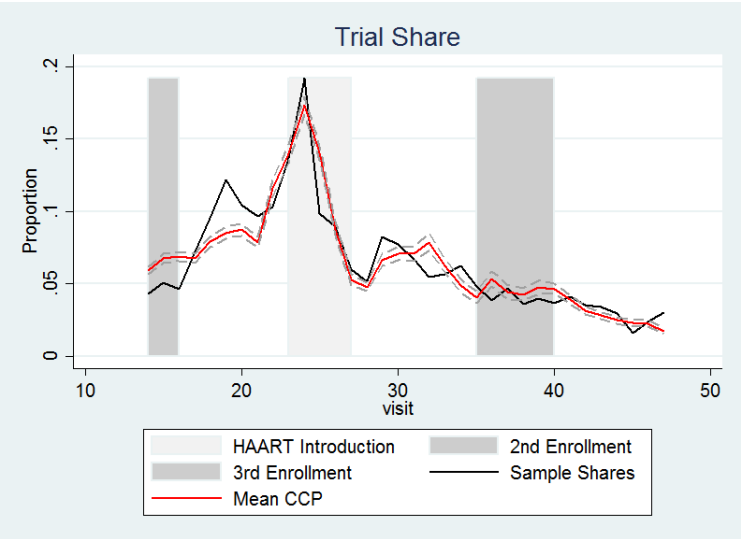

Appendix Figure S3: Average CCPs for Trial

Notes: The figure shows the average estimated conditional choice probability against the share of people choosing the alternative. Dashed lines represent 95% confidence intervals around the predicted CCPs. Three periods of special relevance are highlighted in the Figure: two periods during which enrollment into the sample was undertaken and the period in which products belonging to the HAART class were introduced.

D Results Appendix

D.1 Utility Parameters

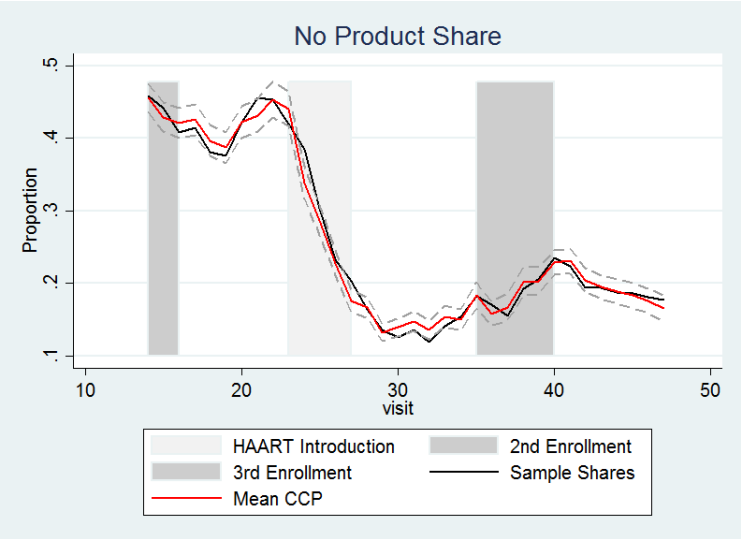

Appendix Figure S4: Average CCPs for No Product

Notes: The figure shows the average estimated conditional choice probability against the share of people choosing the alternative. Dashed lines represent 95% confidence intervals around the predicted CCPs. Three periods of special relevance are highlighted in the Figure: two periods during which enrollment into the sample was undertaken and the period in which products belonging to the HAART class were introduced.

Appendix Table S4: Utility Parameters, y_{it}

parameter	variable	coef.	uncorrected se	se
α_{4w}	$Cluster_{it} \cdot White_i$	-3.546	(0.179)	(0.744)
α_{4b}	$Cluster_{it} \cdot Black_i$	-4.190	(0.190)	(0.762)
α_{4l}	$Cluster_{it} \cdot Hispanic_i$	-3.967	(0.647)	(0.958)
α_{4a}	$Cluster_{it} \cdot Age_{it-1}$	0.043	(0.004)	(0.011)
α_{4h}	$Cluster_{it} \cdot Health_{it-1}/10^3$	-2.021	(0.104)	(0.423)
α_{5w}	$Trial_{it} \cdot White_i$	-1.468	(0.136)	(0.280)
α_{5b}	$Trial_{it} \cdot Black_i$	-2.553	(0.142)	(0.334)
α_{5l}	$Trial_{it} \cdot Hispanic_i$	-1.585	(0.300)	(0.356)
α_{5a}	$Trial_{it} \cdot Age_{it-1}$	0.032	(0.003)	(0.005)
α_{5h}	$Trial_{it} \cdot Health_{it-1}/10^3$	-2.461	(0.078)	(0.203)
α_{6w}	$Stay_{it} \cdot White_i$	0.502	(0.130)	(0.567)
α_{6b}	$Stay_{it} \cdot Black_i$	0.276	(0.145)	(0.613)
α_{6l}	$Stay_{it} \cdot Hispanic_i$	0.707	(0.354)	(0.454)
α_{6a}	$Stay_{it} \cdot Age_{it-1}$	0.009	(0.002)	(0.007)
α_{xp}	$NoAilments_{it} \cdot NoProduct_{it}$	1.019	(0.260)	(1.767)
α_m	$GrossIncome_{it} - OutPocketPay_{it}$	0.057	(0.010)	(0.057)

Notes: Estimation of equation (13). Discount factor $\beta = .95$. $Cluster_{it}$ indicates whether the individual chose one of the three clusters of products available. $Trial_{it}$ indicates whether he chose a trial treatment. $Stay_{it}$ indicates whether he decided to continue using the same treatment he used last period. $NoProduct_{it}$ indicates whether he did not consume a product. $Health_{it-1}$ is defined as the number of white blood cells per cubic millimeter of blood. In parentheses, uncorrected standard errors and corrected standard errors computed using subsampling with 100 subsamples.

D.2 Product Characteristics

Tables S6 and S5 present the estimates of equations (S5) and (S6). Table S5 is an exhaustive list of estimated product characteristics and Table S6 presents the health parameters of both equations.[59] Table S7 presents estimates of the process determining within cluster weights, given by equations (9) and (10).

[59] In separate exercises not shown here we run several versions of equations (S5) and (S6) using lower and higher degree polynomials. A fifth degree polynomial seems to be the best fit for the data.

Appendix Table S5: Product Characteristics

Market Product	Ailments, θ^x coeff	se	Health, θ^h coeff	se	Market Product	Ailments, θ^x coeff	se	Health, θ^h coeff	se
AZT	-0.500	(0.020)	-12.004	(0.736)	ddI , d4T, Nevirapine	0.753	(0.175)	44.240	(3.781)
Interferons (α and/or β), AZT	-0.600	(0.061)	-55.796	(3.102)	ddI , 3TC, Nelfinavir	-0.810	(0.083)	47.816	(6.848)
AL-721 egg lecithin	-0.433	(0.087)	-19.655	(3.917)	ddI , d4T, Efavirenz	-0.626	(0.078)	41.280	(2.772)
AZT, Acyclovir	-0.539	(0.050)	-12.752	(1.670)	3TC, Abacavir, Efavirenz	0.108	(0.047)	53.341	(1.501)
Acyclovir	-0.783	(0.047)	-0.017	(2.678)	AZT, Nevirapine, 3TC, Abacavir	0.038	(0.131)	39.379	(3.369)
AZT, Acyclovir, ddI	-0.851	(0.037)	-16.474	(1.497)	AZT, 3TC, Abacavir, Efavirenz	0.348	(0.080)	78.914	(3.549)
Acyclovir, ddI	-0.348	(0.043)	-4.159	(2.479)	AZT, 3TC, Efavirenz	0.342	(0.079)	43.526	(3.073)
AZT, ddC	-0.439	(0.029)	-5.155	(1.309)	AZT, 3TC, Abacavir	-0.442	(0.078)	54.824	(3.175)
AZT, ddI	-0.571	(0.061)	-16.615	(2.488)	d4T, 3TC, Efavirenz	-0.346	(0.069)	47.978	(3.876)
ddI	-0.375	(0.071)	15.263	(2.587)	Nevirapine, 3TC, Abacavir	-0.470	(0.099)	17.866	(12.148)
AZT, ddC, Acyclovir, ddI	-0.789	(0.115)	-13.351	(7.73)	d4T, 3TC, Kaletra	-0.310	(0.123)	35.611	(5.199)
AZT, ddC, Acyclovir	-0.514	(0.086)	-13.186	(2.168)	3TC, Kaletra, Abacavir	-0.934	(0.124)	51.570	(5.325)
AZT, ddC, ddI	-1.440	(0.047)	-32.700	(1.801)	AZT, 3TC, Kaletra	-0.655	(0.140)	49.838	(3.967)
ddC, Acyclovir	-0.310	(0.093)	2.415	(4.370)	AZT, 3TC, Kaletra, Abacavir	0.298	(0.234)	9.855	(9.404)
ddC	-0.358	(0.084)	-18.630	(3.389)	3TC, Abacavir, Efavirenz, Tenofovir	-0.308	(0.070)	31.845	(3.848)
d4T	-0.717	(0.054)	39.776	(2.210)	AZT, 3TC, Abacavir, Tenofovir	-0.652	(0.074)	19.273	(5.651)
AZT, Acyclovir, 3TC	-0.527	(0.096)	42.267	(3.394)	AZT, 3TC, Kaletra, Tenofovir	-0.552	(0.067)	32.227	(2.681)
AZT, 3TC	0.064	(0.051)	34.398	(1.875)	Nevirapine, 3TC, Tenofovir	-0.258	(0.163)	27.246	(4.619)
Acyclovir, d4T, 3TC	-0.509	(0.100)	33.792	(4.664)	3TC, Kaletra, Tenofovir	-0.092	(0.082)	51.672	(2.709)
AZT, 3TC, Saquinavir	-0.271	(0.052)	38.283	(1.992)	Kaletra, Efavirenz, Tenofovir	-0.966	(0.100)	47.617	(2.684)
d4T, 3TC	-0.104	(0.112)	37.173	(4.070)	3TC, Efavirenz, Tenofovir	-0.011	(0.108)	47.790	(5.468)
AZT, 3TC, Saquinavir, Ritonavir	-0.591	(0.085)	57.776	(10.571)	AZT, 3TC, Kaletra, Abacavir, Tenofovir	-0.738	(0.141)	19.980	(4.226)
AZT, Acyclovir, 3TC, Indinavir	-0.479	(0.056)	63.734	(2.201)	ddI , Kaletra, Tenofovir	-0.276	(0.112)	18.396	(4.015)
Acyclovir, d4T, 3TC, Indinavir	-0.295	(0.108)	78.559	(3.665)	ddI , Efavirenz, Tenofovir	-0.420	(0.117)	2.381	(2.505)
AZT, 3TC, Ritonavir, Indinavir	-0.567	(0.102)	35.032	(6.629)	Abacavir, Efavirenz, Tenofovir	-0.762	(0.140)	39.457	(3.150)
d4T, 3TC, Ritonavir, Indinavir	-0.767	(0.049)	33.510	(3.321)	Kaletra, Abacavir, Tenofovir	-0.820	(0.198)	14.891	(2.601)
d4T, 3TC, Saquinavir, Ritonavir	-0.444	(0.085)	42.631	(5.409)	3TC, Ritonavir, Abacavir, Atazanavir	-0.061	(0.039)	26.850	(1.181)
ddI , d4T, Indinavir	-0.048	(0.137)	32.286	(3.981)	Efavirenz, Tenofovir, Emtricitabine	0.118	(0.082)	54.798	(2.464)
d4T, 3TC, Indinavir	-0.395	(0.096)	53.128	(4.546)	Ritonavir, Efavirenz, Tenofovir, Emtricitabine, Atazanavir	0.306	(0.053)	83.823	(1.706)
AZT, 3TC, Indinavir	-0.075	(0.066)	65.041	(2.809)	3TC, Ritonavir, Abacavir, Tenofovir, Atazanavir	-0.403	(0.163)	38.313	(10.521)
d4T, Nevirapine, 3TC	-0.386	(0.052)	46.846	(2.962)	ddI , Ritonavir, Tenofovir, Atazanavir	0.049	(0.108)	47.800	(2.837)
AZT, Nevirapine, 3TC	0.109	(0.087)	46.275	(4.061)	Ritonavir, Tenofovir, Emtricitabine, Atazanavir	0.138	(0.104)	53.028	(3.940)
AZT, 3TC, Nelfinavir	-0.432	(0.072)	50.776	(3.924)	Nevirapine, Tenofovir, Emtricitabine	-0.205	(0.079)	37.227	(2.303)
ddI , d4T, Nelfinavir	-1.049	(0.060)	57.227	(3.672)	Kaletra, Tenofovir, Emtricitabine	-0.183	(0.093)	46.723	(5.990)
d4T, 3TC, Nelfinavir	-0.881	(0.134)	48.018	(9.588)	Ritonavir, Tenofovir, Emtricitabine, Lexiva	-0.372	(0.116)	30.226	(3.328)
Fringe Mixes									
Isoprinosine, Ribavirin, Interferons (α and/or β)	-1.017	(0.110)	-21.950	(6.644)	Nevirapine, 3TC, Ritonavir, Kaletra, Tenofovir	-1.265	(0.113)	45.683	(4.934)
Interferons (α and/or β), 3TC, Saquinavir, Indinavir, Efavirenz	-0.054	(0.243)	65.353	(5.179)	3TC, Ritonavir, Kaletra, Abacavir, Tenofovir, Atazanavir	-0.465	(0.077)	28.440	(2.687)
Nevirapine, 3TC, Saquinavir, Ritonavir, Indinavir	0.068	(0.134)	6.457	(7.335)	Ritonavir, Tenofovir, Emtricitabine, Atazanavir, Lexiva	-0.612	(0.142)	42.050	(3.579)
Nevirapine, 3TC, Saquinavir, Ritonavir, Nelfinavir	-0.689	(0.156)	30.293	(7.841)	Saquinavir, Ritonavir, Tenofovir, Emtricitabine, Atazanavir	-0.665	(0.120)	31.824	(3.879)
Nevirapine, Saquinavir, Ritonavir, Abacavir, Efavirenz	-1.121	(0.161)	19.278	(4.112)	3TC, Ritonavir, Abacavir, Tenofovir, Atazanavir, Lexiva	-0.210	(0.078)	26.678	(5.890)
Nevirapine, Ritonavir, Nelfinavir, Abacavir, Efavirenz	-0.697	(0.099)	31.044	(4.027)	Saquinavir, Ritonavir, Abacavir, Tenofovir, Emtricitabine	0.072	(0.142)	32.865	(4.856)
Nevirapine, Ritonavir, Kaletra, Abacavir, Efavirenz	-0.410	(0.174)	43.495	(5.757)	3TC, Ritonavir, Tenofovir, Emtricitabine, Raltegravir	0.032	(0.094)	33.352	(2.728)
Nevirapine, 3TC, Nelfinavir, Abacavir, Tenofovir	-0.467	(0.109)	27.893	(3.250)	Ritonavir, Tenofovir, Emtricitabine, Darunavir, Raltegravir	-0.221	(0.067)	47.736	(2.929)

Notes: Product characteristics are estimated as indicators for treatment usage using equations (S5) and (S6). Equation (S5) is a logit model where the independent variable is whether the individual did not suffer ailments during period t. Equation (S6) is a linear model where the independent variable the CD4 count at the end of period t. In parentheses, standard errors computed using subsampling with 100 subsamples. For *Fringe Mixes* we only include the 5 or 6 most used products in the mix.

Appendix Table S6: Health Parameters of Equations (19) and (20)

	Ailments, γ^x		Health, γ^h	
Variables	coeff	se	coeff	se
$CD4_{t-1}$	0.008	(0.0004)	1.152	(0.013)
$CD4_{t-1}^2/10^3$	-0.013	(0.001)	-0.519	(0.043)
$CD4_{t-1}^3/10^7$	0.109	(0.017)	4.375	(0.546)
$CD4_{t-1}^4/10^{10}$	-0.040	(0.010)	-2.016	(0.298)
$CD4_{t-1}^5/10^{14}$	0.054	(0.021)	2.803	(0.546)
Constant	-0.929	(0.038)	-5.874	(1.350)

Notes: Parameters estimated using equations (S5) and (S6). Equation (S5) is a logit model where the independent variable is whether the individual did not suffer ailments during period t. Equation (S6) is a linear model where the independent variable the CD4 count at the end of period t. In parentheses, standard errors computed using subsampling with 100 subsamples.

Appendix Table S7: Within Cluster Weights Function

variable	coef.	se
Ailments Rk	-0.427	(0.124)
Ailments Rk $\times$ *Health Rk*	0.074	(0.020)
Health Rk^2	-0.029	(0.008)
Ailments Rk^2	-0.019	(0.006)
N	-0.509	(0.048)
Health Rk $\times$ N	0.046	(0.009)
Ailments Rk $\times$ N	0.063	(0.010)
Ailments Rk $\times$ *Health Rk* $\times$ N	-0.007	(0.002)
New	-0.352	(0.508)
New $\times$ N	0.027	(0.404)
Constant	0.786	(0.121)

Notes: Parameters estimates from equations (9) and (10) which describe a nonlinear regression of shares within a cluster as a function of product characteristics. *Ailments Rk* stands for the ranking of the ailments component of the characteristics as compared to the other treatments within a cluster; *Health Rk* is defined similarly. N is the number of treatments in the cluster. *New* indicates whether the product just entered the market. In parentheses, standard errors computed using subsampling with 100 subsamples.

D.3 Alternative Regimes

We explore the consequences of the following alternative regimes on the innovation process:

1. *Science regimes*. In these regimes we separate the process of innovation from demand. On the entry margin we transform the centroid to be just a simple average of the characteristics of products currently available in the market—as opposed to the share weighted average in the baseline model (see equation (1)). Because this definition changes the realized innovation values in equation (2), we also estimate new parameters for equation (3), and new distributions of innovations and number of new products, which are associated with the realized innovations directly and through Q_t in equation (5). However, since we need a trial participation path to determine the distribution of the number of new products, we set the trial participation path at the average from the baseline simulated paths. In this way we keep that part of the comparison constant relative to the baseline. On the exit margin we exogenously drop products from the market based on their quality. For this we follow one of two procedures as explained below:

   ↪ *Frontier*. Any product that is not in the technological frontier is dropped from the market.

   ↪ *Thick frontier*. The exit rate path is set at the average from the baseline simulated paths. This exit rate determines the number of products n_t to be dropped. Then, given n_t, the products that are dropped from the market are chosen from inside the *thick frontier*, which is formed with the following algorithm. Given the set of products available at t ($\mathcal{P}_t$), define the set $I_1 \equiv \mathcal{P}_t$. Then starting from iteration $k = 1$, follow the steps:

   (a) Build the kth frontier as:

   $$F_k = \{\theta \in I_k : \nexists \theta' \in I_k \text{ s.t. } \theta' \geq \theta\} \tag{S30}$$

   (b) Define the new inside set as:

   $$I_{k+1} = I_k \setminus F_k \tag{S31}$$

   (c) If $\#I_{k+1} > n_t$, set $k = k + 1$ and go back to Step (a).

   (d) If $\#I_{k+1} \leq n_t$, drop all products in I_{k+1} and randomly drop $(n_t - \#I_{k+1})$ from F_k. Stop.

2. *Random choice.* In this regime we study the evolution of product quality when the process of innovation responds to demand but demand is random. Random demand neutralizes the dependence of the technological path on the preferences and characteristics of consumers without changing the nature of the process on the supply side. In practical terms, we maintain the definition for the centroid as a share-weighted average (see equation (1)). As a consequence, the distributions of innovations and number of new products remain unchanged. Moreover, we aim to avoid a spurious effect of arbitrary aggregate shares on the process of innovation. Therefore, instead of allocating arbitrary unconditional probabilities for each choice (f.i. $1/G$ for a choice set of size G), we fix the unconditional probabilities of the random choice regime to match the unconditional shares in the baseline.

D.4 One-Time Planners

We study the policy decisions of one-time planners that internalize the externalities in our model but who act only for one period. These planners do not change the regime characteristics, which allows us to use the conditional choice probabilities obtained in estimation to construct continuation values induced by the planners' choices. We do not solve the planners' problems analytically. Instead, we search over all possible policy rules and choose the one that maximizes average utility. This process is computationally burdensome because we need to simulate forward for each policy rule, and there can be millions depending on how the state space is sliced. To address this issue we rely on two devices. First, we either constrain the information the planner has or her choice set. Second, we precompute a set of continuation values and match them with planner rules using the state induced by the planner's rule. Both procedures are further explained below.

D.4.1 Constrained Planners

Information constraint. The first planner we consider is constrained in the amount of information she has. This planner can only observe whether a person's health at the beginning of the period is high or low and whether the person decided to consume a market treatment last period (either by staying on his previous market product or by trying a cluster). Hence, the planner's policy rules can be based only on four different categories. Using her limited information, the planner can send individuals to any of the six alternatives available, insofar as they are feasible for the individual.[60] We nest the baseline individually optimal solution by

[60]The six choices are: three clusters, staying in previous market product, trial treatment and no treatment. Recall that only people who consumed a market treatment last period get to stay in a market product.

adding one alternative to the planner's choice set: we let the planner rely on the individual's information by letting the individual choose his optimal treatment given his information.[61] Hence, the planner has 7 alternatives and can base her policy rules on 4 categories. Since only two of the four categories can stay in a previous market treatment, this amounts to $7^2 \times 6^2 = 1,764$ policy rules. An example of a policy rule is presented in Table S8.

Appendix Table S8: Example of an Information-Constrained Planner's Policy Rule

Category		Alternatives						
Health status	Product $t-1$	Cluster 1	Cluster 2	Cluster 3	Trial	Stay	No product	Free
high	yes							x
high	no		x					
low	yes				x			
low	no						x	

Notes: *Product* $t-1$ column indicates whether individuals in this category consumed a market product in $t-1$ either by staying on their previous market product or by trying a cluster. *Free* column indicates that the planner relies on the individual's information by letting the individual choose his individually optimal treatment given his information.

Choice constraint. The second planner we consider is constrained in the choices she has. This planner has full information of the individual's state but can only do one of two things: send the person to trials or let the person choose what is individually optimal among the remaining five alternatives. The policy rules for this planner are levels of trial participation. Since this planner if fully informed she orders people in terms of who gains/losses the most from going to trial relative to their second best alternative. Then, for a given policy rule, she assigns to trials the individuals who like the alternative the most until she reaches the policy rule's level of trial participation. This planner problem nests the individually optimal solution when the planner's policy rule is exactly the individually optimal share of trial participation. For policy rules below the individually optimal share of trial participation the planner incurs a welfare costs by preventing people from joining a trial who wanted to join. For policy rules above the individually optimal share of trial participation she incurs a welfare costs by forcing people to join a trial who did not want to join. The welfare gains come from the externality via experimentation in clinical trials. The number of policy rules to be evaluated depend on how finely we discretize the increments in trial participation, which is itself constrained by the number of people in the sample.[62] We evaluate policies in increments of 0.5 percent points, which amounts to 202 policy rules.

[61] This requires that the individual does not act strategically in respond to his position as a free chooser relative to people who are assigned choices. Otherwise, the conditional choice probabilities would have to be adjusted to account for this strategic behavior.

[62] The minimum increment size possible is $1/N$.

D.4.2 Continuation Values

To facilitate the solution of the planner problems above we avoid forward simulating a number of continuation values for every rule and averaging over the simulated paths. Instead, we implement the following steps to reduce computational burden:

1. We create a collection of 500 continuation value vectors (each row in a vector is an individual) computed for as many next-period states. We denote this collection $\mathcal{A}$, where each component $m \in \mathcal{A}$ contains a continuation value B^m and a next-period associated state Ξ^m.

2. For each rule n in a given planner problem, we compute each individual's current payoff and their future state, as well as the implied aggregated next-period state Ξ^n_{t+1}.

3. We match rule n with the continuation value in $\mathcal{A}$ corresponding to the next-period state in $\mathcal{A}$ that is closest to Ξ^n_{t+1}. In other words, we match rule n with the continuation value $B^{m^\star}$ given by the indicator m that solves:

$$m^\star = \min_{k,\ \Xi^k \in \mathcal{A}} ||\Xi^n_{t+1} - \Xi^k|| \tag{S32}$$

   We use a measure of Euclidean distance yielding from discretizing the aggregate states Ξ^n_{t+1} and Ξ^k into vectors with $R = 196$ components. We scale each component r in every aggregate state vector to be between zero and one by dividing over the largest component r across all vectors.

4. We repeat steps 2 and 3 200 times for every rule n and average over repetitions.

Smoothing. When finding the solution to the problem of the planner who is choosing the optimal level of trials participation, our method to speed up computation generates some noise around the mapping from planner rules into average consumer life-time utility (see Figure S5). Hence, we use a local polynomial to smooth the mapping in an interval starting at the individually optimal share and going 15 percent points above it (from .09 to .24 in Figure S5). This produces Figure 14 and the results associated with it in Table 11.

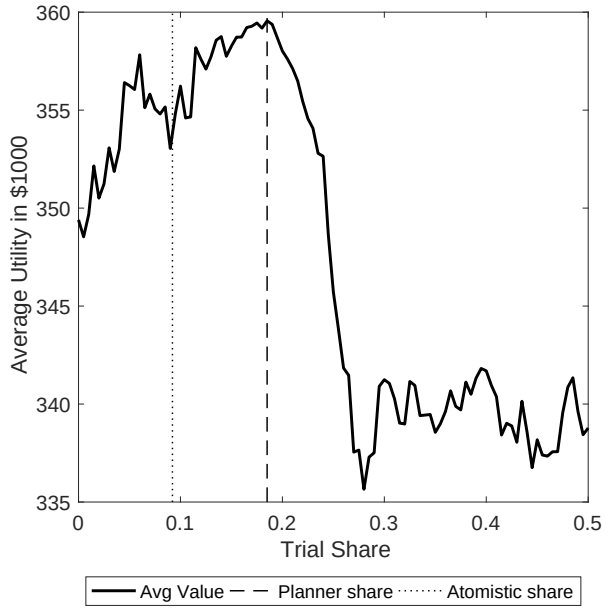

Appendix Figure S5: Optimal Assignment to Clinical Trials

Notes: The solid line represents the average welfare generated by a constrained planner in the second semester of 1996 that sends individuals to trials according to a rank based on how individuals like going to trials relative to their individually optimal choice. The dotted line indicates the individually optimal share. The dashed line indicates the planner's optimal share.