

The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D*

Zarek Brot-Goldberg[†]

Timothy Layton[‡]

Boris Vabson[§]

Adelina Yanyue Wang[¶]

January 5, 2021

Abstract

We show that default rules in the low-income portion of Medicare Part D have large and persistent effects on enrollment and drug utilization. We develop a framework for optimal default policy and show that a key unknown parameter is the elasticity of active choice propensity with respect to the value of the default assignment. Using randomly-assigned defaults, we estimate an attention elasticity close to zero, even when beneficiaries face substantial losses in drug consumption. Variation in active choice rates appears to be largely random. Beneficiaries are likely to benefit from paternalistic or "smart" defaults rather than be hurt by them.

*We thank Jason Abaluck, Stefano DellaVigna, Tal Gross, Martin Hackmann, Ben Handel, Jon Kolstad, Brigitte Madrian, Michael McWilliams, Joe Newhouse, Mark Shepard, Anna Sinaiko, Justin Sydnor, Dmitry Taubinsky, and seminar participants at the 2019 American Society of Health Economists meeting, the BU/Harvard/MIT Health Economics Seminar, Georgetown McDonough, Harvard Medical School, the 2019 Midwest IO Fest, the Yale Economics Junior Workshop, and the 2021 ASSA Meetings for helpful comments. We gratefully acknowledge support from the National Institute on Aging under award number P01AG005842 (P29-D26) and award number P01AG032952 and from the Agency for Healthcare Research and Quality (the Agency for Healthcare Research and Quality (K01-HS25786-01)). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or AHRQ. All errors are our own.

[†]University of Chicago. Email: zarek@uchicago.edu

[‡]Harvard University and NBER. Email: layton@hcp.med.harvard.edu

[§]Harvard University. Email: vabson@hcp.med.harvard.edu

[¶]NBER. Email: adywang@alumni.stanford.edu

Across a variety of domains, default rules appear to have large effects on individual behavior. When choosing their retirement savings account contribution rates, individuals overwhelmingly receive the default option (Madrian and Shea 2001, Choi et al. 2002). Opt-out default rules increase enrollment in organ donation (Johnson and Goldstein 2003), electricity pricing schemes (Fowlie et al. 2017), and vaccination (Chapman et al. 2010) relative to opt-in rules. In health insurance markets, year-after-year consumers tend to remain enrolled in their prior plans (typically the default), despite large financial losses from doing so (Handel 2013, Polyakova 2016).

Despite all of this evidence that consumers follow defaults, we know relatively little about *why* they do so. The theoretical and empirical literatures have offered competing hypotheses, including costs of decision-making, procrastination, and inattention (Bernheim and Taubinsky 2018). But we have little evidence as to which of these models actually drives behavior. Understanding which of these explanations is responsible for the well-documented phenomenon of consumers overwhelmingly following defaults is not just an academic curiosity; models of how attention is formed are critical to designing optimal default rules, and ‘nudge’ policies more generally. Indeed, the lack of understanding of which explanation is the correct one has led to a conflict in the literature between those advocating paternalistic defaults that match defaults to consumer preferences (Thaler and Sunstein 2003, 2008, Handel and Kolstad 2015) and those suggesting that “shocking” defaults that are purposefully *mis*-matched to consumer preferences in order to induce consumers to make active choices might be the optimal default choice (Choi et al. 2003, Carroll et al. 2009, Beshears et al. forthcoming). If passivity is driven by a boundedly-rational bias such as procrastination, consumers may settle for suboptimal defaults rather than opt for their most-preferred choice, and a shocking nudge may spur them to action. Conversely, if consumers are simply forgetful or randomly inattentive, they may follow such a default into harm’s way. There is currently little guidance for policymakers and market designers regarding which direction their default designs should take.

In this paper, we attempt to provide evidence that could be used to provide such guidance for a setting where the question of the optimal default is an urgent one affecting millions of Americans every year: Health insurance. The dominant form of health insurance provision in the United States is one in which beneficiaries are given a choice of heterogeneous insurance plans in a regulated marketplace, both in private provision by employers and in public provision (Gruber 2017). These marketplaces are characterized by a surprising diversity of default rules, many of which have unclear theoretical or practical bases. For example, defaults range from re-enrollment in one’s prior plan in most employer-sponsored insurance markets, to ‘smart’ defaults in the Maine Medicare Part D market for dual-eligibles, to randomly-assigned defaults in many state Medicaid programs and in the national Medicare Part D Low-Income Subsidy Program. Recently, more extreme defaults have been proposed, with the Trump administration proposing to change the default for continuing enrollees in the Health Insurance Marketplaces created by the Affordable Care Act from automatic re-enrollment to automatic re-enrollment *without subsidies* (Federal Register 2020).

We study this question in two parts. First, we provide novel evidence that, in the domain of health insurance, consumers do indeed overwhelmingly follow defaults, both in terms of their initial choices and in terms of their continuing choices, and that this default-following has (negative) consequences for those consumers. Second, we provide a framework for evaluating the differences between models that imply default-following behavior, both in the sense of their testable implications, which we evaluate empirically,

and in their implications for normative policymaking. Our analysis focuses on the Medicare Part D Prescription Drug Insurance Program in the United States, one of the largest regulated health insurance markets in the world. Throughout, we focus on the Low-Income Subsidy (LIS) segment of the program, which provides subsidized coverage to low-income program beneficiaries.

We start by using two unique natural experiments in the LIS program to provide novel evidence that consumers follow defaults. We first provide evidence that *new* LIS beneficiaries follow defaults. To do so, we leverage the fact that a few months before a new LIS beneficiary initially qualifies for Medicare, an algorithm randomly selects a default prescription drug plan for her from a set of pre-specified plans that qualify for LIS auto-assignment. If the beneficiary does not opt out of this default and make an active choice prior to the start of Medicare enrollment, she is enrolled in the randomly-chosen default plan. Our data allows us to observe the initial randomized assignment *and* whether the beneficiary was enrolled through the default assignment rule. We document that *only 16% of new LIS beneficiaries opt out of their randomly-selected default plan and actively choose a different plan* during initial enrollment, with the remaining 84% ultimately being automatically enrolled in the randomly-selected plan. Of those who were enrolled in their default assignment, nearly two-thirds did not make an active choice to opt out of their default after five years, a striking degree of choice passivity. Indeed, in any given year, *two-thirds of all LIS beneficiaries are enrolled in a plan that they were originally enrolled in by default (random) assignment*. Further, we find that even beneficiaries for whom the potential value from opting out appears to be the highest rarely do so.

While this exercise shows that defaults matter in initial assignment, we cannot immediately conclude that the persistence in default plan enrollment is due to the fact that the typical default is to remain in one's plan over time, or due to the presence of real costs of switching plans, such as filling out paperwork to reauthorize prescriptions, adjusting to a new formulary that may not cover the drugs a patient is currently taking, etc. When the default is to remain in one's plan, these explanations cannot be separately identified. Using a second natural experiment, we next provide evidence that continuing LIS beneficiaries also follow defaults, *even when that default differs from their prior plan*. To do so, we leverage unique exogenous variation in the default faced by these returning beneficiaries. Each year, Part D plans must submit a bid to the Medicare regulator that determines the premium charged to non-LIS beneficiaries. After these bids are submitted, the LIS subsidy is calculated as an average of those plans' bids. For beneficiaries enrolled in benchmark plans in year t that end up with a bid for year $t + 1$ *below* the $t + 1$ subsidy level, the default is to remain in their incumbent plan. However, for beneficiaries in plans whose $t + 1$ bid falls *above* the $t + 1$ subsidy, the default is to be automatically switched to a different, randomly-chosen default plan. Because there are many plans bidding close to the subsidy level in any given year, and because the subsidy is unknown ex-ante, this gives rise to a regression discontinuity design, where we compare the plan switching rate of beneficiaries whose incumbent plan's bid was just below the subsidy to those whose bid was just above the subsidy.

We estimate that the effect of the default changing from re-enrollment in the incumbent plan to re-assignment to a different plan raises the extent to which beneficiaries change plans from nearly 0% to 97%. This dramatic rise in the switching rate is driven entirely by reassignment through the default mechanism rather than active choice triggered by the change in the default. The impact of switching is long-lasting: Roughly 99% of those who do not initially opt out of the default of reassignment remain in their new plan for at least two years post-reassignment. Most beneficiaries do not appear to be willing to pay as little as \$2

per month to remain in their incumbent plans. These results provide strong evidence that consumers follow defaults in this market, no matter what that default may be, and largely reject the hypothesis that the high levels of choice persistence in this program are explained by consumers actively choosing to remain in their plans in order to avoid real switching costs.

Ultimately, these results, that new LIS enrollees overwhelmingly stick with their randomly assigned default plans for years and that returning LIS enrollees passively follow randomly-assigned defaults rather than actively choosing to remain in their prior plans, show that defaults are a powerful tool in the regulator's toolbox. Given that power, this begs the question of how regulators should optimally set default rules. Should default assignments be close to beneficiary preferences and minimize opt-outs (Thaler and Sunstein 2003, 2008, Handel and Kolstad 2015)? Or should default assignments be purposely *mis*-aligned with beneficiary preferences in order to 'shock' them into making an active choice (Choi et al. 2003, Carroll et al. 2009, Beshears et al. forthcoming)? In order to answer this question, we develop a general framework of discrete choice under cognitive costs that nests the majority of models of consumer choice typically used in this literature. Importantly, our framework allows for the possibility that either paternalistic *or* "shocking" defaults could be optimal. We use the framework to show that in many cases knowing the exact model of consumer choice is not necessary for determining optimal policy. Instead, we show that, in general, optimal default policy involves a fundamental trade-off: A decrease in the value of the default decreases welfare for inframarginally passive beneficiaries but increases welfare for beneficiaries who are newly incentivized to make an active choice and move to their best option. We show that, rather than determining the underlying model of consumer choice, determining the optimal default primarily requires knowledge of a critical parameter: The extent to which beneficiaries are in fact triggered to make active choices by defaults that are increasingly harmful to them. If the relationship between the value of the default and the portion of consumers making active choices is non-negative (i.e. increasingly harmful defaults do *not* induce more active choice), then policymakers face no trade-off and paternalistic defaults will unambiguously improve consumer welfare.

Much of the literature uses parametric models of attention that implicitly or explicitly impose assumptions about this parameter. We use our unique setting to estimate it directly. First, we show that different possible default rules induce different 'stakes' for beneficiaries, so that the default-following behavior we observe is not a consequence of a rationally-inattentive agent giving little attention to a decision that is inconsequential. Specifically, we show that beneficiaries who are exogenously switched from the default of remaining in their prior plan to re-assignment to a randomly-selected default plan experience significant losses of drug consumption. Event study estimates show that reassignment to a randomly-selected benchmark plan results in a 6.4% drop in drug spending on average, with substantial declines in usage of high-value drugs used to treat chronic conditions. Significant declines persist when we standardize prices across plans, suggesting that these declines reflect actual declines in drug *consumption* rather than differential drug prices across new and incumbent plans. We further leverage the randomization of defaults to beneficiaries to estimate that the consequences of re-assignment versus remaining in one's prior plan are much more extreme for those beneficiaries assigned to plans whose formularies are a poor match for the set of drugs taken by the beneficiary. Plans that are low-value according to a proxy measure of the value to beneficiaries (the portion of the beneficiary's drugs covered by the plan's formulary) decrease drug spending

by 12.6% (compared to a default of not switching plans), while plans that are higher-value decrease drug spending by only 4.3%. This gap again persists when we standardize drug prices across plans and when we zoom in on a basket of high-value drugs that physicians deem especially important for beneficiary health, indicating real *ex-post* differences in the value of these plans to beneficiaries.

We then leverage this variation in the value (as proxied by formulary fit) of the various plans from which beneficiary default options are randomly selected to estimate how the propensity to make an active choice is related to the value of the default, the parameter identified by our framework as being critical to determining the optimal default. We show that the effects of the value of the default on the probability of making an active choice are extremely small: Beneficiaries assigned to the highest-value plans for them are only 2.3 percentage points more likely to make an active choice than beneficiaries assigned to the lowest-value plans for them. To put that number in context, no matter what default plan the beneficiary is assigned, the share of beneficiaries making an active choice is fewer than 10 percent. This is true despite our finding that beneficiaries assigned to the low-value plans experience a loss of drug consumption *three times* the size of the loss experienced by the beneficiaries assigned to the high-value plans. This result persists at least *two years after assignment*, and it also holds for the subset of beneficiaries for whom the difference in fit between the high- and low-value plans is largest. Our results thus provide strong evidence that for the range of defaults assigned in our setting (which involve substantial differences in beneficiary consumption), there is little relationship between the value of the default and the propensity to make an active choice. To the extent that this result generalizes to other settings, the key implication is that proposals to provide paternalistic defaults (Thaler and Sunstein 2003, Handel and Kolstad 2015) are much more likely to improve beneficiary welfare than proposals attempting to use defaults as an incentive to induce active decision-making. This is due to the fact that any welfare gains from the very few beneficiaries induced by a worse default to make active choices are highly likely to be vastly outweighed by welfare losses incurred by the majority of beneficiaries who passively follow that worse default.

Given that the value of the default does not appear to explain variation in active choice, in our final section, hoping to shed light on the ‘correct’ underlying model of choice operating in this setting, we ask what *does* seem to explain this variation. To do so, we use a different natural experiment to decompose inattention into an individual-specific component (i.e., some are attentive and others are not) and a transitory component (i.e., a given beneficiary is sometimes attentive and sometimes not). Identifying this requires us to observe a sequence of active choice events. Observing a positive covariance of active choices for the same beneficiary across time would imply that the individual-specific component would have a greater weight. We leverage the fact that, when a beneficiary’s incumbent plan *exits*, the default is auto-assignment for *all* plan enrollees,¹ and ask whether those beneficiaries who had previously actively enrolled in their plan are more likely to make an active choice again. We find that they are, with an active choice rate of 25.6% compared to 7.7% for the previously-passive. A simple calibrated structural model suggests that two-thirds of the variation in (latent) attention is due to transitory shocks to attention, with the remainder explained by individual heterogeneity. Further, only around 3% of the variation (9% of the individual heterogeneity) can be explained by characteristics that are observable in standard Medicare data, such as demographics or

¹In contrast, the default change used in our regression discontinuity approach only affects beneficiaries who had been previously auto-assigned to their incumbent plans.

health status. Ultimately, these results suggest that much of the variation in active choice appears ‘random’ and there is thus little scope for targeted ‘nudges’ that are directed at particularly attentive or inattentive ‘types’ of beneficiaries. Instead, broad-based default policies may be the most effective.

Our paper makes important contributions to several literatures. First, we contribute to a robust literature documenting mistakes and behavioral biases in health insurance choice. While there is substantial prior work on inertia in health insurance choice, papers in this literature typically impose specific, often parametric, assumptions about when, why, and how beneficiaries are affected by biases. For example, [Handel \(2013\)](#), [Ericson \(2014\)](#), and [Polyakova \(2016\)](#) implicitly impose what we call a ‘rational inattention’ model in which beneficiaries exactly trade off the value of their default plan against the value of their best alternative minus the cost of attention, and effectively bound the cost of attention from below by the (parametric) forgone value of following the default.² In contrast, we are able to identify the extent and consequences of inattention without making any parametric assumptions about beneficiary preferences, and our results suggest that the models leaned upon in this prior work assume too much ‘rationality’ in how beneficiaries respond to the value of their default. Moreover, this lack of rationality may solve the puzzle of the extremely large cognitive/switching cost estimates that the literature has produced.³

We also extend the general literature on the extent to which default rules drive behavior in various settings. Our result that default rules matter in health insurance is perhaps unsurprising given similar effects in retirement savings ([Choi et al. 2002, 2004](#), [Beshears et al. 2009](#)), among a large literature whose review is beyond the scope of this paper. However, in our setting, compared to many of these others, the cost of a mismatched default is realized immediately for a beneficiary who discovers that her preferred drug is uncovered, compared to a saver who may not realize for years that the ‘path of least resistance’ leaves them with too little or too much retirement savings. Despite this, we still find large and extremely persistent default effects. Our analysis is closest to that of [Blumenstock et al. \(2018\)](#) and [Andersen et al. \(forthcoming\)](#), who similarly use field data to distinguish between different models of choice passivity. In contrast to prior work, including theirs, we are able to directly estimate distortionary follow-on effects of default misalignment on consumption, rather than infer these effects from ex ante mismatch.⁴ We also supplement a small literature on optimal default policy that has largely focused on the 401(k) retirement savings context ([Choi et al. 2003](#), [Thaler and Sunstein 2003](#), [Carroll et al. 2009](#), [Beshears et al. forthcoming](#)). That work has compared optimal defaults across many standard models of default-following behavior ([Bernheim et al. 2015](#)) under various normative assumptions ([Goldin and Reck 2020](#)). Our approach provides a simple characterization of the optimal default across a wide variety of plausible models, in terms of the normative quantity employed by these prior papers (which we refer to as the extent of ‘mistakes on the margin’), as well as a positive quantity, the elasticity of active choice with respect to the value of the default, a parameter that can be

²Alternative strategies in the literature include imposing the assumption that active choice is *only* responsive to the value of the default and not that of other alternatives ([Abaluck and Adams-Prassl 2020](#)); that default-following is driven by year-by-year *changes* in plan features and not their levels ([Ho et al. 2017](#), [Heiss et al. 2016](#)); or that beneficiaries perfectly observe some features of their plan options but not others ([Brown and Jeon 2020](#)). These strategies also often define implicit rational or boundedly-rational models of attention formation.

³As a modal example, [Handel \(2013\)](#) estimates yearly switching costs of, on average, \$2,032, which is a substantial fraction of the estimated value of insurance. Models of boundedly-rational attention formation would need lower cognitive/switching costs to rationalize the same data.

⁴One exception, [Fowlie et al. \(2017\)](#), undertake a similar analysis in the electricity pricing context.

estimated from data. Importantly, our approach does not require a policymaker to take a strong stance on the specific underlying behavioral model.

Finally, we join a growing literature on models of endogenous attention. Models of consumer theory under ‘(boundedly-)rational inattention’ seek to understand how the extent of biases of market participants may be endogenous to the structure of the market (Gabaix 2014, 2019, Matějka and McKay 2015). We show that there is an approach with minimal assumptions that allows an economist to measure the ‘rationality’ of attention by using variation in the ‘stakes’ of the decision. Our approach echoes that of recent experiments done in lab and survey settings that elicit costs of attention, such as Caplin et al. (forthcoming), Morrison and Taubinsky (2020), and Bronchetti et al. (2020). Our work is closest to recent work estimating attention elasticities to stakes in field settings by Beshears et al. (forthcoming) in prescription drug delivery and Brown and Jeon (2020) in the unsubsidized portion of Medicare Part D. Our work is distinct from theirs in that we have quasi-experimental variation in the decision stakes, rather than purely observational variation. In contrast to much of this new body work, we find that attention is (perhaps surprisingly) largely *insensitive* to the stakes in the market. Our results may more closely reflect real-world attention elasticities compared to lab settings which inherently focus the subject’s attention.

The paper proceeds as follows: In Section 1, we describe Medicare Part D, the Low Income Subsidy program and its associated default rules, and the data we use to examine default effects. In Section 2 we leverage random assignment of default plans to provide evidence that new LIS beneficiaries follow defaults. In Section 3 we use our regression discontinuity design to separate out default effects from other sources of choice persistence and to show that continuing LIS beneficiaries follow defaults. In Section 4 we develop a framework for understanding the behavior we observe and develop a theory of optimal defaults. In Section 5 we leverage the natural experiments to estimate the ‘stakes’ of following the defaults and use those estimates to infer relative value of the defaults. In Section 6 we estimate a key parameter from the framework, the opt-out response to decision stakes. Section 7 uses an additional natural experiment around plan exits to decompose the sources of variation in attention. Section 8 concludes.

1 Institutional Setting and Data

1.1 Medicare Part D and the Low Income Subsidy Program

Medicare Part D was introduced in 2006 by the Medicare Modernization Act of 2003 to fill a coverage gap in Medicare, which had previously lacked any outpatient prescription drug benefit. Under Part D, drug coverage is fully outsourced to privately contracted insurers providing coverage on the government’s behalf. The Medicare program organizes a centralized market in which beneficiaries may select from one of these private plans. In 2016, Medicare Part D covered about 41 million beneficiaries and accounted for about \$94 billion of annual expenditures, of which 86% was paid for by federal and state governments and the remainder by individual beneficiaries through premium payments and cost-sharing.

Part D plans are required to adopt a standardized benefit design, in terms of their financial as well as non-financial features. Specifically, plans follow a standardized cost-sharing structure and are required to cover at least two drugs in each of 148 therapeutic classes. Nonetheless, plans are given some scope for

differentiation, in terms of the specific drugs that they cover on their formulary, the specific cost-sharing levels they assign to different formulary tiers, and non-price rationing mechanisms (such as authorization restrictions) that they institute for covered drugs. Meanwhile, consumers enjoy choice among a substantial variety of Part D plans, being free to choose any one of the many Part D plan offered in their service region. In 2016, the average beneficiary had access to 26 stand-alone Part D plans.

While the Part D program is heavily subsidized for all beneficiaries, those with financial need are granted additional subsidies through the low-income subsidy program (LIS), which offers supplemental drug premium and cost-sharing support. Around 30% of Medicare beneficiaries participate in the LIS program. So-called ‘dual-eligibles’ (those who simultaneously qualify for both Medicare and their state’s Medicaid program) are automatically enrolled in the LIS program when they qualify for Medicare, as are beneficiaries in the Medicare Savings Program. Others not automatically eligible for LIS, but who meet income and asset eligibility criteria, can qualify by applying directly.

The generosity of LIS subsidies varies by group, with non-dual LIS recipients receiving partial premium and cost-sharing support, while the dual-eligibles on which our study focuses receive much more generous subsidies. Specifically, dual-eligibles receive a subsidized reduction in their plan premium payments up to the ‘benchmark’ amount, which is set at the average of plans’ premium bids in that region for that year. This premium subsidy covers the full premium amount for a subset of plans, meaning that dual-eligibles enrolling in those plans (otherwise known as ‘benchmark plans’) would have no premium responsibility. They do not get to keep the difference if the subsidy exceeds the premium of their plan. Those dual-eligibles opting to enroll in non-benchmark plans do not get their premiums fully covered by the subsidy, and consequently would be left with some out-of-pocket responsibility equal to the difference between the plan’s premium and the subsidy. Multiple benchmark plans are available in a service region; beneficiaries have access to between two and sixteen, with 92% of beneficiaries having at least 5 benchmark plans to choose from (see Appendix Figure A9).

Dual-eligibles in the LIS program additionally receive substantial cost-sharing subsidies: They have all plan deductibles completely waived, under both benchmark and non-benchmark plans. In addition, dual-eligibles in both benchmark and non-benchmark plans are shielded from each plan’s drug-specific copayment and coinsurance schedule, which can otherwise be a significant source of out-of-pocket costs for Medicare beneficiaries. Instead, dual-eligibles face their own custom copayment schedule: In 2020, they were charged a copayment of \$1.30 for all formulary-covered generic drugs and \$3.90 for all formulary-covered branded drugs, though in most cases these nominal copayments are not actually collected. In addition to substantially reducing out-of-pocket costs, this policy also makes plans effectively uniform in their financial characteristics for dual-eligible beneficiaries, nullifying any variation in cost-sharing determined by the plan sponsor.

Given that dual-eligibles’ Part D out-of-pocket expenses are uniform for covered drugs, plans most materially differ in the set of drugs covered by their formularies (the set of drugs that they offer coverage for), along with the utilization management requirements imposed on those drugs that are covered. Formulary exclusion of a drug means a beneficiary would have to pay the full sticker price of that drug out-of-pocket with no subsidization if they opt to purchase the drug. Because regulation requires plans to cover at least two drugs in each therapeutic class, beneficiaries will generally have another covered drug (potentially

of substantially different match quality for their particular needs) that they can switch to as a substitute. Alternatively, they can switch to a plan that covers their chosen drug.

This heterogeneity is important. There is substantial cross-plan variation in the set of drugs covered by plans' formularies. To take the popular anti-cholesterol drug Lipitor as an example, of the nine benchmark plans in New York in 2009, six plans covered the drug on their formulary while three did not. Further, even among the six plans that did cover the drug, two required prior authorization for beneficiaries to obtain coverage, while four did not. As such, if a LIS beneficiary in New York wanted to fill a prescription for Lipitor but was on one of the three plans that did not cover it, they had the option to either pay for its cost 100% out-of-pocket, substitute to another drug, or switch plans entirely. Given that different anti-cholesterol medications (as well as many other classes of important drugs) are not perfect substitutes for one another, these cross-plan differences in drug coverage can be meaningful for beneficiaries and carry real consequences, which we explore in Section 5.

1.1.1 LIS Part D Default Rules

What makes the LIS program unusual, and particularly suitable for our study, is the default rule it uses to assign beneficiaries who do not make an active plan choice. The LIS default rule is designed to help low-income beneficiaries avoid plans with positive premiums while simultaneously not favoring any particular zero-premium plan. This results in each LIS beneficiary being randomly assigned a benchmark plan as their beneficiary-specific default. We now briefly sketch a more detailed picture of how this default works in practice.

When an individual who already qualifies for a state Medicaid program turns 65, and thus qualifies for Medicare as well, she is automatically enrolled in the LIS program. A few months before they are automatically enrolled in the LIS program, beneficiaries are assigned to a plan that is chosen at random from among the set of stand-alone benchmark plans in their service region. Benchmark plans are fully subsidized by Medicare and thus have zero premiums for beneficiaries. LIS beneficiaries are informed of their assigned plan and can opt out of this default at any point, including before their initial Part D enrollment is effectuated, by actively selecting from among any of the plans available in their service region, including non-benchmark stand-alone Part D plans as well as combined Medicare Advantage-Part D plans; neither of the latter plan types would be in scope for random assignment under the default mechanism. Unlike unsubsidized beneficiaries, LIS beneficiaries can switch plans at the start of *any month*, not just during the open enrollment period at the end of the year. If opting to actively choose a plan, LIS beneficiaries can select the plan that best fits their existing prescription drug basket using publicly accessible information on the set of prescription drugs covered by each plan in their choice set. Furthermore, CMS's official plan selection portal ('Plan Finder') offers decision support tools that can automate the prescription-drug-basket-to-plan matching process for beneficiaries, by soliciting information on the drugs they currently take and then automatically suggesting the plans that best fit their needs.

After initial enrollment, continuing LIS beneficiaries face a different set of default rules. In most markets for insurance products, including the unsubsidized component of Part D, the default is generally for the enrollee to remain in the plan they were enrolled in during the prior year. This is generally true in the LIS

program as well, with one major exception: When plans lose their ‘benchmark’ status and are thus no longer free to LIS beneficiaries, the default for beneficiaries who had previously been assigned to that plan switches from remaining in the prior plan to automatic enrollment in a randomly-selected (from among benchmark plans in the local market) beneficiary-specific default plan. In October of year $t - 1$ these beneficiaries receive notice that their default plan for year t is *not* their current (year $t - 1$) plan but instead a plan that has been randomly chosen for them from among the set of year t benchmark plans in their market. They are told that if they do not opt out of this new plan before January, they will be automatically switched on January 1. They can also choose to opt out anytime after January 1 and revert back to their prior plan or enroll in any other available plan. This change in the default does *not* apply to beneficiaries who had actively chosen the plan that lost benchmark status.

The motivation for this change in the default is to avoid defaulting a beneficiary into a plan charging a premium. Plans lose benchmark status when they bid above the subsidy for the coming year. If the plan to which a beneficiary was previously assigned loses benchmark status, default re-enrollment would thus leave the beneficiary paying for a plan they did not opt into (because they were initially assigned to the plan via the auto-assignment mechanism). The default is designed to avoid that situation, nudging beneficiaries into (free) new benchmark plans rather than nudging their continued enrollment in newly non-benchmark plans with positive premiums.

A final set of default rules applies to LIS beneficiaries in Part D plans that cease to exist entirely. Those beneficiaries enrolled in an exiting plan are randomly assigned default plans from the set of available benchmark plans (similar to new Medicare beneficiaries). These defaults apply to both those who had been auto-assigned to the exiting plan as well as those who had actively elected it. This is distinct from the unsubsidized population, whose default upon plan exit is disenrollment from the Part D program altogether.

So far, we have described the default mechanism as randomly assigning default plans to beneficiaries. While this is true in a general sense, there are notable caveats and exceptions to this rule, which we account for in our analysis. First, assignment is undertaken in a ‘stratified random’ method: Rather than true random assignment to benchmark plans, beneficiaries are first randomized across insurers who offer benchmark plans, and then randomized across plans within their assigned insurer. This is to discourage the potential practice of introducing duplicate plans to artificially increase the receipt of assigned beneficiaries. Second, in the case in which beneficiaries are re-assigned (i.e. if their plan loses benchmark status), if their incumbent plan’s insurer offers another benchmark plan in their service region, they are routed to that plan rather than randomly assigned. Finally, auto-assignment is not random in Maine for dual-eligible beneficiaries, as Maine is the only state to have implemented a ‘smart default’ policy. Rather than assign default plans at random, Maine uses prior data on beneficiary drug consumption to assign them the plan that best covers their prescriptions. For these latter two cases, we exclude all affected beneficiaries from our subsequent analyses.

1.2 Data

We make use of several administrative datasets from the Centers for Medicare and Medicaid Services (CMS). These data contain information on beneficiary program enrollment status, medical utilization, and

prescription drug utilization across the Medicare and Medicaid programs. The Medicaid and Medicare data are linkable at a beneficiary-level, allowing us to observe health conditions and drug utilization prior to initial Medicare and Part D enrollment at 65 for some beneficiaries. The Medicare data is nationwide in scope and extends from 2007 to 2015, tracking drug utilization for all Medicare beneficiaries and medical utilization for all beneficiaries outside of Medicare Advantage. The Medicaid data only covers the states of New York and Texas for the 2004-2010 time period. However, it tracks medical as well as drug utilization for all Medicaid beneficiaries, including those enrolled in private Medicaid Managed Care plans, across these two states.

Beneficiary Demographics, Enrollment, and Choice Status. We obtain information on beneficiary demographic characteristics and plan as well as program enrollment from the Medicare Beneficiary Summary File and the Medicaid Personal Summary File. For Medicaid and Medicare beneficiaries, these files provide important demographic information such as age, gender, and geographic location, including the Part D plan region to which individuals belonged that year. The Medicare files additionally track enrollment status at a person-month level for different Medicare programs, including Part A (the hospital benefit), Part B (outpatient), Part C (Medicare Advantage), and Part D (prescription drugs). These Medicare files also track dual Medicare-Medicaid enrollment status at a person-month level, along with whether individuals are enrolled in the LIS program and whether they qualify for the full LIS subsidy (note that all dual-eligibles automatically qualify for the fully-subsidized version of the LIS). Separately, the Medicaid files track Medicaid program enrollment at a person-month level, along with beneficiary characteristics. The beneficiary IDs in the Medicaid and Medicare files are standardized to a common format, enabling us to track any given individual across both data sources, and, for example, follow an individual as they transition from Medicaid-only to dual-coverage at age 65.

We combine this data with a newly released plan election type file. This file covers all Part D enrollment spells from 2007-2015, and for each spell tracks whether enrollment was initiated through active choice or default auto-assignment. Importantly, in addition to listing the plan a beneficiary was enrolled in during each month, the file also includes the default plan that was assigned to the beneficiary, even if the beneficiary opted out of that default. This allows us to observe the *assigned* plan as well as the *realized* plan for each beneficiary.

Plan Characteristics Data. We obtain information on plan characteristics from publicly available CMS datasets, which cover all Part D plans offered during our sample period. These data track the set of plans offered in each Part D plan region. For each plan in each year it was offered, we are able to observe the monthly premium that the plan charged (both the Medicare-paid portion and the beneficiary-paid portion); the plan's benchmark status; and the entirety of the plan's formulary, including both which drugs were covered and, for covered drugs, whether those drugs faced some kind of authorization restriction, such as prior authorization, step therapy, or quantity limits.

Outpatient Prescription Drug Data. We track outpatient prescription drug usage for a random 20% sample of Part D enrollees using claims-level information from both the Medicare and Medicaid programs,

based off the Part D Event files and MAX RX files, respectively. Each claim represents an event where a beneficiary filled a single prescription of a given drug. For each claim, we observe the specific drug prescribed and filled (at the NDC code level), the quantity/days supply for the fill, as well as the date when the prescription was filled, and the cost charged to the beneficiary and to the Medicare program.

Analytic Sample. Throughout the paper, we utilize a variety of subsamples of LIS beneficiaries to exploit multiple natural experiments. For all of these analyses, we employ a set of common restrictions for any subsample we construct. We restrict only to dual-eligibles (those simultaneously enrolled in Medicaid and Medicare, and thus receiving full LIS benefits) over 65 enrolled in Medicare Parts A, B, and D, and not enrolled in Medicare Advantage. We generally sample at the beneficiary-year level and require these restrictions to be true for every month in a year in which we include a beneficiary in our sample. We describe summary statistics for all beneficiaries captured in this analytic sample in the first column of Table 1. The population in question is relatively aged and has substantial medical use, both for prescription drugs and for other medical services.

2 Do New LIS Beneficiaries Follow Defaults?

In this section we leverage our first natural experiment to show that LIS beneficiaries follow defaults. We start by focusing on beneficiaries who are eligible for Medicaid (and thus the LIS program) and who have *newly* qualified for Medicare coverage by turning 65. Unlike other groups who must actively choose to take up Medicare Part D, individuals dually eligible for Medicaid and Medicare are automatically enrolled in Part D at age 65. In the months before their 65th birthday they receive a letter in the mail informing them that they will be enrolled in Part D in the month of their 65th birthday and that they will be enrolled in a randomly-chosen benchmark plan (listed in the letter) unless they opt to actively choose a different Part D plan. In other words, these beneficiaries are randomly assigned defaults. By determining the extent to which beneficiaries ultimately enroll in and remain enrolled in their randomly assigned default plans, we can determine the extent to which defaults matter for these individuals. If defaults do not matter, beneficiaries assigned different defaults should be equally likely to enroll in a given plan. If defaults do matter, then beneficiaries assigned a particular default plan should be much more likely to enroll in and remain enrolled in that plan relative to beneficiaries assigned other default plans. We report summary statistics for our sample, restricted to those who turned 65 between 2007 and 2015, in the second column of Table 1. Given the relative youth and health of these individuals compared to our broader sample, these are individuals who should have greater cognitive ability than the average LIS beneficiary.

Figure 1a presents the portion of new LIS beneficiaries who actively opt out of their default, either through actively choosing a plan before automatic enrollment takes effect (in month zero), or by actively switching out of their plan post-enrollment (months one and on). We plot the cumulative share of beneficiaries who have opted out as a function of the number of months since their initial enrollment in Medicare. We restrict only to those for whom we observe five years of enrollment data following their initial qualification for Medicare, i.e., those who aged into Medicare between the years of 2007 and 2010. We define active choosers here as those making independent, explicit choices of plans. We exclude *re-assignments* to

new plans through the default reassignment mechanisms we detailed in Section 1.1.1. We average across all cohorts who enter Medicare during our sample window.

Opt-out is rare: Only 16% of LIS beneficiaries newly enrolled in Medicare make an active plan choice between being notified of their randomly-chosen default plan and their date of enrollment in Part D. The other 84% of beneficiaries enroll in their randomly assigned default plan. This clearly shows that defaults matter in this setting. While these initial default rates are strikingly high, they could potentially be explained by the fact that beneficiaries can switch plans each month, causing them to not feel pressure to make a plan choice immediately, especially if they have limited drug needs. However, even after five years of enrollment, only 45% of all beneficiaries have *ever* made an active opt-out decision. This 45% is comprised of the 16% who opted out before they turned 65, and an additional 29% (approximately one-third of those who initially enrolled in their randomly-assigned default plan) who opted out at a later point in the first five years of Medicare enrollment. This persistence is striking given that the initial default was enrollment in a *randomly-selected* plan.

In order to examine whether these strikingly low rates of active choice persist for various sub-groups of beneficiaries, such as those with high drug needs or in particularly poor health, we construct a secondary sample made up of a subset of LIS beneficiaries for whom we observe Medicaid claims prior to enrollment in Medicare at age 65. This sample allows us to construct *ex-ante* measures of beneficiary characteristics like health and drug needs using data from periods prior to Part D enrollment, eliminating the possibility that these measures capture plan effects on health and drug use rather than underlying beneficiary characteristics. This sample consists of beneficiaries living in New York and Texas who aged into the Medicare program during the years 2007-2010. We report summary statistics for this sample in the third column of Table 1. In Figure 1b, we replicate Figure 1a for this subsample. While the subsample is generally less likely to make an active opt-out decision (with only 8% making an active opt-out decision before initial enrollment), they are otherwise qualitatively similar on other observable characteristics, except that they are slightly sicker.

Active Choice by Health Status. In Figure 2a we stratify beneficiaries in this sample by health status (as measured by the Elixhauser comorbidity index) at age 64, prior to their enrollment in Medicare. At 65, differences in active choice across the groups are relatively small, with the healthiest (no chronic conditions) making an active choice 3% of the time, while 7% of the sickest (at least four chronic conditions) do so. Over time, the groups diverge somewhat, but after five years only 45% of the sickest and around 33% of the healthiest beneficiaries have made an active choice, again indicating an extraordinarily high level of choice passivity in this setting. Appendix Figure A1 presents a similar analysis where we stratify beneficiaries by the number of unique drugs they took at age 64. Again, even among the beneficiaries using many drugs (13+) at age 64, only 50% make an active choice by the end of their 5th year of enrollment in Part D.

Active Choice by Value of Choice. While measures of health status are suggestive of beneficiaries' benefits from choice, they are *not* direct measures of that benefit. A relatively healthy beneficiary taking a single, rarely-covered drug will face poor coverage for her care in expectation under the default mechanism despite the fact that a plan exists to cover all her drugs. By contrast, a very sick beneficiary may not face much benefit from choosing between drug plans if she takes only drugs that are frequently covered by formularies

in benchmark plans.

In the spirit of this example, we construct a more direct proxy for the value of choice to an individual. We first identify the set of drugs the beneficiaries in our sample were using at age 64 using Medicaid claims. Then, for every plan in the beneficiary's initial choice set at the time she entered Medicare, we construct a measure of 'fit.' We define beneficiary-plan fit as the share of the drugs the beneficiary was taking at age 64 that are covered by the formulary of the plan. We discuss this measure of fit in depth in Appendix B, including showing (in Appendix Figure A12) that there is significant variation in fit across plans, ranging from 100% coverage to 0%, with substantial mass in between.

For beneficiaries, a plan with poor fit will require them to either change their drug consumption behavior (substitute from uncovered drugs to covered drugs) or pay the full list price. Thus, the value of choice should be higher for beneficiaries assigned default plans of relatively poorer fit. For each beneficiary, we rank the benchmark plans available to them by their person-plan 'fit' measure. Two beneficiaries in the same market can have vastly different personal rankings based on their prescriptions. In Appendix B we show (in Appendix Figure A13) that, for a given person, the difference between the average and the best benchmark plan in terms of 'fit' can be a difference of as large as 40% or more of their age 64 basket of drugs.

We divide our sample into quintiles based on each beneficiary's own personal ranking of their initial randomly-selected default plan. A beneficiary assigned to the benchmark plan in their market that was the worst fit (among the set of benchmark plans) for their drug basket is in the fifth quintile, while the beneficiary assigned to the plan which best matched their prescriptions is in the first quintile. For each of these groups, we plot their cumulative active choice propensity over time in Figure 2b. Strikingly, the propensity to make an active choice remains stubbornly low even for those with the most to gain from doing so. Even for the group with the worst possible assignment, less than 10% opt out of the initial randomly-chosen default plan before their enrollment is effectuated ($t = 0$), less than 25% have made an active choice after 2 years, and less than 45% have made an active choice after 5 years in Part D. Further, there seems to be no monotonic relationship between the fit of a beneficiary's randomly-assigned default plan and her propensity to make an active choice. We return to this fact in Section 6 where we establish the robustness of this result and discuss its implications for the optimal default in this setting.

Summary. These stylized facts provide clear evidence that default assignments matter for beneficiary plan enrollment in this setting. None of the subgroups of beneficiaries most likely to benefit from choice as measured by health status, use of drugs, or the relative formulary fit of their assigned default plan have initial rates of active choice above 10%. Further, all of these groups have rates of active choice below 50% *even after 5 full years in the program*. This level of choice persistence is also striking, especially when recognizing that the initial assignment was to a randomly-selected plan. Our results here also rule out the standard neoclassical explanation that the choice persistence we observe is caused by persistent preferences, i.e. that beneficiaries make an initial active choice based on their preferences and they stick with it over time because their initial choice continues to be aligned with their preferences (Keane 1997). Since the initial assignment was not a choice, but instead randomly assigned, the initial plan of enrollment is orthogonal to beneficiary preferences over plan features, and thus choice persistence does not reflect a specific preference

for that plan.⁵ Instead, this pattern must reflect some kind of state-dependence in plan choice.

3 Do Continuing LIS Beneficiaries Follow Defaults?

In this section, we leverage a second natural experiment to show that defaults matter not only for new beneficiaries but also for continuing beneficiaries. In Section 2 we showed that the majority of beneficiaries do not opt out of their default during the first five years of enrollment in Part D, suggesting that continuing beneficiaries, like new beneficiaries, also follow defaults. However, default-following is not the only possible explanation for the choice persistence we document in Section 2. After the initial assignment the default is to remain in the initially assigned plan. Thus, choice persistence can either be due to passive default-following *or* beneficiaries actively avoiding switching plans (i.e. state-dependent preferences). In a model of state-dependent preferences, all plans other than the beneficiary’s incumbent plan get a utility penalty in consideration (Handel 2013, Ericson 2014, Polyakova 2016). Under such a model, this utility penalty, or “switching cost,” explains choice persistence above and beyond neoclassical considerations. Any alternative to the incumbent plan needs to provide additional value of at least the switching cost, or else the beneficiary will not be willing to pay the penalty to switch to it. Such a penalty might include the cost of adapting to a new formulary, which could require finding substitutes for drugs that were covered by the beneficiary’s incumbent plan but not by her new plan, and/or reauthorizing drugs for a new insurer. While choice passivity must explain our initial assignment results, the persistence of the initial assignments across time could reflect either continued passivity *or* an avoidance of switching costs.

It is important to differentiate between choice persistence due to switching costs versus choice persistence due to passivity/default-following because the two explanations for choice persistence have very different implications for default rule design. If switching costs explain choice persistence, altering the default will have little effect—beneficiaries will still actively opt to remain in their plans in order to avoid the switching cost. But if switching costs are not the driving factor behind choice persistence, defaults can be powerful tools for shifting beneficiary choices. In this section, we provide a clean test for distinguishing between these two explanations for persistence among continuing beneficiaries.

3.1 Default Effects vs. Switching Costs

Distinguishing between passivity and switching cost explanations is impossible when the default option is for the beneficiary to re-enroll in her incumbent plan (Handel 2013). Both sets of mechanisms will lead to the same observable outcome: Beneficiaries will remain in their plans. Instead, in order to adjudicate between the two, we need a setting where consumers who are roughly identical, including in terms of their cost of switching plans, face different defaults. Such a setting would allow us to observe whether consumers

⁵In practice, some beneficiaries will be randomly assigned to the best option for them. This cannot rationalize the results we find here. First, some of this heterogeneity would be picked up in our analysis of opt-out rates by the fit of the randomly-assigned default plan, and the fact that we do not find any heterogeneity by fit suggests limits to such a possibility. Second, the probability of being assigned to one’s own personal most-preferred (benchmark) plan is $\frac{1}{N}$, where N is the number of benchmark plans. Even with $N = 2$, the smallest number of benchmark plans ever available in any region, the probability is $\frac{1}{2}$, which is less than the share of assigned beneficiaries who fail to opt out even after five years.

are just as likely to follow the default when it is some alternative option versus when it is to re-enroll in their incumbent plan.⁶

The unusual default rules for continuing LIS beneficiaries, described in detail in Section 1, provide exactly this type of setting. For beneficiaries who were previously auto-assigned to their plan, the typical default is to remain in their incumbent plan. This default remains in force *unless the plan loses its benchmark status by setting a year t premium above the year t subsidy*. If a plan that was a benchmark plan in year $t - 1$ sets its premium in year t above the year t subsidy, the default for its incumbent enrollees switches from remaining in the plan to being auto-enrolled in a *different*, randomly-selected year t benchmark plan. This setting provides us with a clear test of whether persistence is driven by defaults: If so, a change in the default rules will induce beneficiaries to switch plans through auto-assignment. In contrast, if persistence is driven by switching costs, beneficiaries whose default changes to a different randomly-chosen benchmark plan will actively choose to remain in their incumbent plan in order to avoid incurring the switching cost.

In order to deal with the possibility that beneficiaries enrolled in plans that lose benchmark status may differ from beneficiaries enrolled in plans that do not lose benchmark status, we implement this test via a regression-discontinuity (RD) design. Specifically, we compare year-to-year plan switching rates for beneficiaries whose incumbent plans set premiums just below the subsidy, and therefore retained their benchmark status, versus beneficiaries whose incumbent plans set premiums just above the subsidy, losing their benchmark status.

Research Design. For this analysis, we create a subsample of all beneficiaries who were enrolled in a benchmark plan in December of year $t - 1$ and had been enrolled in that plan through the auto-assignment mechanism.⁷ We describe this subsample in the fourth column of Table 1. This subsample is slightly younger than the general LIS population, and has slightly less medical utilization, but is otherwise qualitatively similar.

Our research design is based on comparing the switching behavior of beneficiaries whose incumbent plan's year t bid was just under the year t subsidy level to those whose plans bid just above the subsidy level. This design is valid under the assumption that whether a plan bids just above or just below the subsidy level is orthogonal to the latent propensity of beneficiaries to make an active choice. One case for this orthogonality is that premium bids are, in general, quasi-randomly distributed around the subsidy. One might worry that plans would bid just under the subsidy level to avoid shedding any default-following customers. Our strategy leverages the fact that the subsidy is set *after* plans have already set their premiums, as the subsidy is set as a (weighted) average of plan premium bids within each Part D region. While plans would prefer to be just below the subsidy and therefore receive the maximum subsidy amount without losing

⁶One candidate is to compare initial choices (where the default is random assignment) and later choices (where the default is to remain in the incumbent plan) as we did in Section 2. This setting is not ideal for testing the switching cost hypothesis, however, as the beneficiaries initially entering the Medicare program are not identical to the beneficiaries making choices after having been enrolled in the program for some period of time.

⁷We restrict to beneficiaries who had been enrolled in the plan through the auto-assignment mechanism because these are the beneficiaries whose default changes when the plan bids above the benchmark. As we discuss below, the default for beneficiaries who actively chose the plan is to remain in the plan, even if the plan bids above the benchmark. We also restrict our sample to beneficiaries whose incumbent plans continued to exist in the following year and additionally restrict to Part D plans classified as 'basic' through the following year to ensure comparability between benchmark and non-benchmark plans; all benchmark plans have the basic coverage classification, while non-benchmark plans are allowed to offer an 'enhanced' plan type if they so choose.

any default-following customers, we assume that they cannot *perfectly* foresee what the subsidy will be, and their position relative to the benchmark cutoff thus represents forecasting error. One might worry that some plans are more likely to make forecast errors than others. In practice, plans frequently lose benchmark status from one year to the next, with, for example, 44% of all 2008 benchmark plans losing that status in 2009.

To assess our assumption of quasi-randomness, in Figure 3, we plot a histogram of the incumbent plans' premiums (relative to their regional subsidy level) in the following year, in the spirit of the density test of McCrary (2008). We note two important features of this plot. First, in our sample, more plans set premiums below the subsidy cutoff than above. Approximately 68% of beneficiaries are enrolled in a plan that set a premium below the cutoff (while 12% are included in 'de minimis' plans which bid above the cutoff but were allowed to revised their bids down, as we describe below). This is a consequence of our requirement that our sample only contains plans that were benchmark plans in the prior year and thus all had premiums below the subsidy. A rational forecast of their future prices would be for the mean to remain under the subsidy amount with some variation.⁸ Importantly, however, the distribution of plan premiums appears to be relatively smooth across the discontinuity.

Second, we have highlighted plans in the \$0-\$2 bin range as potentially being 'de minimis' plans. This refers to a policy that CMS adopted in 2011 (and in some areas in earlier years). For plans whose bids were in this range just above the subsidy, the plan could, ex post, choose to move their bid down to the subsidy and retain their existing auto-assigned beneficiaries rather than having them get re-assigned. They could not, however, receive new auto-assigned beneficiaries. In practice, nearly every plan exercised this option. Given the take-up was not 100%, however, we drop all plans that bid in this range from 2011 onwards.⁹

We also perform a number of balance tests to show that a rich set of beneficiary characteristics, including prior medical and drug spending, are smooth and comparable across the discontinuity. In the five panels of Figure 4, we plot, as a function of the incumbent plan premium bid relative to the subsidy, various beneficiary characteristics, including average age, share female, average drug spending in the prior year, average non-drug medical spending in the prior year, and the average Elixhauser Comorbidity Index. Regression estimates for differences in these characteristics across the discontinuity are found in Appendix Table A1. None of these tests shows any change in these characteristics around the discontinuity.

Regression Discontinuity Results. To implement our test, for each beneficiary in our sample we measure whether or not the beneficiary retained the same plan between December of year $t - 1$ and January of year t (the default for plans with year t bids below the year t benchmark). In Figure 5, we plot the average switching rate for whole-dollar bins of the difference between each incumbent plan's premium in year t and the regional subsidy in that year. All 'de minimis' plans are dropped. Our results are striking: the average probability of switching among beneficiaries on the left side of the cutoff, where the default is to remain in the incumbent plan, is less than 1%. In contrast, the average probability for beneficiaries on the

⁸In Figure A2 we present a scatterplot relating premiums in $t - 1$ to premiums in t for *all* plans that continued to exist in consecutive years. The conditional mean of the premium in t relative to the cutoff is approximately a linear function of the same quantity in $t - 1$.

⁹In Appendix Figures A3 and A4 we recreate this histogram from the pre- and post-de minimis time periods. The pre-de minimis time period still has a few de minimis plans, owing to the fact that CMS experimented with de minimis rules in a very small number of markets during 2007 and 2008. We drop plans that bid in the de minimis range in relevant markets during this period.

right side of the cutoff, where the default is to be reassigned, is approximately 97%. We interpret this as strong evidence of the effects of defaults—when the default changes from re-enrollment in one’s prior plan to automatic switching to a different randomly-selected plan, beneficiaries do not actively choose to remain in their incumbent plan to avoid incurring a switching cost.

On their own, these results do not conclusively demonstrate that plan switches are coming from the change in the default, as some individuals may be actively switching to a plan that is not their randomly-chosen default plan. As we discuss in Section 4, many common models of default-effects predict that active choice can be ‘triggered’ by a change in the default. Appendix Figure A5 replicates Figure 5 but breaks out the switching rates into active switches and passive switches, as observed in the data. Beneficiaries are, on average, 2.4 percentage points more likely to make an active choice to switch when their incumbent plan’s premium in year t is on the right side of the year t cutoff. This increase in active choice is small compared to those who switch via the default auto-assignment mechanism, indicating the change in the default had only a small effect on the propensity to make an active choice.¹⁰ We discuss the implications of this result for the optimal default in this setting in Sections 4 and 6 below.

Next, we estimate the change in the probability of switching at the discontinuity via regression. We follow Lee and Card (2008) and estimate a global conditional expectation function to approximate the switching rate for beneficiaries whose plans’ premiums were above the subsidy level in the counterfactual scenario where default rules do not change. To do so, we estimate the following linear probability model for an individual i in time t , previously enrolled in plan j in region r , restricting to beneficiaries enrolled in plans with benchmark status in year $t - 1$ whose bids for year t fall within \$6 of the year t benchmark:

$$\begin{aligned} \Pr(\text{Switch Plans})_{it} &= \beta \cdot 1\{Bid_{j(i)t} - Benchmark_{r(i)t} > 0\} \\ &+ \gamma^- \cdot (Bid_{j(i)t} - Benchmark_{r(i)t}) \cdot 1\{Bid_{j(i)t} - Benchmark_{r(i)t} \leq 0\} \\ &+ \gamma^+ \cdot (Bid_{j(i)t} - Benchmark_{r(i)t}) \cdot 1\{Bid_{j(i)t} - Benchmark_{r(i)t} > 0\} \\ &+ \delta X_{it} + \epsilon_{it} \end{aligned}$$

Our parameter of interest is β , which we interpret as the effect of the change in defaults. We present the coefficient estimates from our regression in Table 2.¹¹ In columns 2-4, we divide the effects into active switching, active switching to a benchmark plan, and passive default reassignment (which, by construction, never occurs for beneficiaries in plans below the cutoff). Our regression results echo those from our figures: Incumbent plan benchmark status loss results in a substantial share of beneficiaries switching plans (96%), primarily through default reassignment (93.6%), and only a trivially small portion (2.4%) of beneficiaries are “shocked” by the change in the default into making an active choice.

¹⁰In Figure A5, we also plot the share of beneficiaries who make an active choice and switch to a different *non-benchmark* plan. This is a trivial proportion of switches (0.4pp).

¹¹Lee and Card suggest clustering standard errors at the level of discrete running variable value. We follow the spirit of this inference approach by clustering standard errors at the incumbent plan-by-year level.

Alternative Explanations. One concern with interpreting β as purely the effect of a change in defaults is that remaining in a plan whose premium is above the benchmark requires the enrollee to pay a positive premium equal to the difference between the plan’s year t bid and the year t benchmark. If beneficiaries in our sample are extremely price-sensitive, either due to preferences or to liquidity constraints, our results might instead reflect a rational response to that sensitivity rather than the effect of a default. We present several pieces of evidence that suggest that price sensitivity is not responsible for the dramatic increase in switching.

First, we estimate that γ^+ is approximately zero, so the probability of switching is not increasing in the additional premium that beneficiaries have to pay, i.e. there is no ‘dose response’ of demand to price. Rationalizing this with price-sensitivity requires a model in which beneficiaries’ price sensitivity only takes on extreme values—either zero or a very high value. However, as prior research has shown (Shampanier et al. 2007), changing prices from \$0 to \$1 may have larger effects than going from \$1 to \$2 due to the salience of zero or to extreme liquidity constraints. To explore this, we expand our sample to include beneficiaries in the same plans who had initially *actively* chosen those plans. For these beneficiaries, the default is *always* for the beneficiary to remain in their incumbent plan, whether or not it remains a benchmark plan. In other words, unlike the beneficiaries who were auto-assigned to the plan (those included in the RD analyses above), these beneficiaries experience the same default no matter whether their plan bids above or below the benchmark. In Appendix Figure A6 we compare our plot from Figure 5 to a similar exercise for beneficiaries who were “prior choosers.” We do see an increase in the switching rate at the discontinuity for these prior choosers despite the lack of a change in the default: The switching rate increases discontinuously by about 20% at the cutoff. This does provide evidence that beneficiaries do exhibit sensitivity to positive premiums. However, this sensitivity does not account for anywhere near the full effect seen by the prior non-choosers, at most accounting for 20-25% of the effect for that group, assuming price sensitivity is similar for the prior choosers and the beneficiaries in our main sample.

Another concern is that beneficiaries can switch plans at any time, so accidental assignment to the default mechanism may be reversed the following month when beneficiaries realize they have been switched. In this case, our approach would overestimate the effects of defaults on (long-run) plan switching by only looking at enrollment in January. We address this by performing the same analysis but instead comparing enrollment in December of year $t - 1$ to enrollment in *December* of year t . We present the associated plot in Appendix Figure A7. Our estimates are nearly identical to those that use January enrollment, as only 0.5% of reassigned beneficiaries switch back to their previous plan during year t . Further, in Appendix C we extend the analysis through 24 months after the implementation of the new default and show that these default effects persist throughout that entire 2 year period.¹²

Finally, one might argue that our result that beneficiaries follow defaults does not generalize. The change in the default applied only to beneficiaries who had previously been auto-assigned to the plan. Thus, our results are based on the behavior of beneficiaries who already revealed themselves as the most passive consumers by being auto-assigned to the plan in the first place. To respond to this argument, we first point

¹²The barrier to doing this two-year follow-up in our main specification is that many beneficiaries in both our treatment and control groups are in plans that lose benchmark status the following year, and thus may face a second reassignment within the analysis window. Interpreting treatment effects in such a context is challenging. In Appendix C we restrict to only those beneficiaries whose plans do not lose benchmark status in the following year. However, this restriction reduces our sample size substantially.

out that, at any given point in time, two-thirds of LIS beneficiaries are in a plan they were auto-assigned to. This group thus represents the majority of the LIS program. Further, in Section 7, we use a different natural experiment that includes prior active choosers and show that they too follow the default the majority of the time, indicating that this choice passivity is a general phenomenon among the LIS population.

Summary. Our interpretation of the sum of these results is that we can reject the hypothesis that switching costs drive much of the year-to-year persistence of individual enrollment in a given plan. Instead, defaults seem to be a powerful force in determining the plan choices of continuing LIS beneficiaries, as with new beneficiaries. To clarify, our results do not suggest that real ramifications of switching do not exist; they merely reject the hypothesis that beneficiary choices respond to those costs when they arise. Our results show that beneficiaries are not even willing to pay \$3 per month to remain in their incumbent plan. Moreover, while default-following behavior at a beneficiary’s initial enrollment might reflect difficulty in searching across plans in a new choice environment, such a barrier is not at play for continuing enrollees, where beneficiaries can pay a small premium to simply remain in their incumbent plan and know that their existing set of drugs will continue to be covered.

4 Theoretical Framework

In Sections 2 and 3 we showed that defaults have powerful effects on the choices of LIS beneficiaries. LIS beneficiaries follow defaults not only when the default is to remain in one’s prior plan but also when the default is re-assignment to a randomly-selected plan. In this section we develop a theoretical framework for considering *why* beneficiaries follow defaults and the implications for optimal defaults in this setting, and possibly other settings as well.

We start by providing a general framework for discrete choice with costly attention that nests various models of active choice from the behavioral literature. We show how welfare consequences of default effects differ across various positive and normative models of behavior. We then use this framework to (1) show how different models imply different optimal policies and (2) show that the question of whether the optimal default is paternalistic in the spirit of [Thaler and Sunstein \(2003\)](#) or ‘shocking’ in order to trigger beneficiaries to make an active choice in the spirit of [Carroll et al. \(2009\)](#) depends on a single critical parameter which we attempt to estimate in Section 6.

4.1 Framework

In our framework, agents are given a set of plans to choose from, indexed by $j \in \mathcal{J}$. Each plan provides some utility v^j to the agent who enrolls in it. We denote $v^* = \max_{j \in \mathcal{J}} v^j$ to represent the value of the highest-utility option available to the agent. Agents, when making an active choice of plan, do so to maximize their utility. This choice is to pick $j^* = \arg \max_{j \in \mathcal{J}} v^j$, and they receive a payoff v^* to doing so.

Given this simple framework, we can now nest within it a set of potential models that serve to explain default effects. Because we have already determined that the choices of most beneficiaries in our setting are affected by defaults, we do not consider models where there is no role for defaults and agents always make

active choices. These models include the standard neoclassical economic model, where choices perfectly reflect agent preferences and any choice persistence reflects stable preferences, and the state-dependent preferences and switching cost models, where choices are persistent because beneficiaries are actively avoiding a real switching cost. Under both of these models, there is no role for the default because beneficiaries make active choices in every period. The evidence we presented in Sections 2 and 3 is clearly inconsistent with these models, as we have shown that defaults have powerful effects.

Instead, we restrict our attention to models that allow for the possibility that agents fail to make an active choice.¹³ If the agent successfully makes an active choice, events proceed as above: The agent chooses j^* and receives payoff v^* .¹⁴ However, if the agent fails to make an active choice, she will be given a default option d , with an associated value v^d . Making an active choice requires paying a decision cost c . A decision rule determines whether the agent makes an active choice or not: The agent makes an active choice if and only if $A(v^*, v^d, c; \theta) \geq 0$, where $A(\cdot)$ is an arbitrary function of the value of the default option v^d , the payoff from making an active choice (given by the value of the most-preferred option, v^*), and the decision cost c , given some other factors θ . The decision cost c is a single parameter that includes the various frictions involved with making an active choice. In our setting, it represents the costs, both opportunity costs and psychic costs, of searching for information about the available options; it reflects the hassle costs of filling out the paperwork required to switch; and it reflects the cognitive cost of effort in comparing the options. Under this framework, empirical default effects arise when $A(\cdot) < 0$ for a nontrivial share of agents.

The function A we have defined is generic, and we can think about the empirical literature as estimating parameters of models nested within this framework, with parametric assumptions made on A . To build intuition, for the remainder of this section we use a convenient parameterization of $A(v^*, v^d, c; \theta)$, though none of the conceptual results we develop require it:

$$A(v^*, v^d, c; \theta) = \beta(v^* - v^d) - c + k \quad (1)$$

Under the parameterization, the agent trades off the benefit of making an active choice, $(v^* - v^d)$, against the cost of doing so, c , putting some relative weight β on the benefits versus the costs. They also face random welfare-irrelevant (in the sense of [Bernheim and Rangel \(2009\)](#)) transitory shocks to attention, k , drawn from a distribution $F(\theta)$. We assume that v , c , and k are denominated in the same units.

In Figure 6a we plot the probability of making an active choice (which we refer to as $a(\cdot)$), as a function of the gap in value between the default plan option and the option that would be chosen under active choice, $v^* - v^d$. This can be thought of as the ‘stakes’ at play in making an active choice. When the gap is larger, an agent has more to lose by not making a choice. We consider three specific parameterizations of Equation 1 that are exemplary of common models used in the literature and discuss each in turn. While we refer to these models as representing ‘attention,’ what we describe as ‘paying attention’ can also be thought of as consideration, or overcoming hassle costs, or any other costly decision effort.

¹³One class of models we do not consider are models in which defaults act as ‘suggestions’ to enrollees. Given that beneficiary defaults are randomly selected and described as such, a reasonable enrollee could not infer their assignment as a suggested optimum from Medicare.

¹⁴We assume that an active chooser considers all non-default options. Some recent work has estimated that consumers in insurance markets do not consider all available options ([Coughlin 2019](#)). We make this assumption for tractability, and a model of more limited consideration conditional on active choice would not change the qualitative predictions and welfare considerations here.

Rational Inattention. First, in dark blue, we plot choice as a function of the value gap for a fully rational agent for whom $\beta = 1$, $k = 0$ always, and $c > 0$ is a constant. In such a model, the agent *never* makes an active choice if $v^* - v^d - c < 0$, and *always* does so if $v^* - v^d - c \geq 0$. Since the welfare gains to making an active choice are $v^* - v^d - c$, we describe this parameterization as representing a ‘rationally attentive’ chooser, who only makes an active choice when the expected benefit is greater than the cognitive cost. This is the model used by [Handel \(2013\)](#), and is the general workhorse model in the literature on active choice in health insurance and other markets for contracts. Models of rational search ([Hortaçsu and Syverson 2004](#)) and rational inattention ([Gabaix 2014](#), [Matějka and McKay 2015](#), [Brown and Jeon 2020](#)) follow a similar active choice rule, in which agents rationally trade off costs and benefits, albeit replacing v^* with the expected v^j received as a result of searching or being attentive given the agent’s prior information and c with the expected total cognitive cost they will incur.

Boundedly-Rational Inattention. Second, in green, we plot the active choice probability as a function of the value gap for an agent for whom, again, c is constant and $k = 0$, but now, $0 < \beta < 1$. The difference between this model and the rational inattention model is obvious: The curves have the same shape, but the new agent always makes an active choice if $\beta(v^* - v^d) - c \geq 0$, i.e. if $v^* - v^d \geq \frac{c}{\beta}$. This agent is ‘boundedly-rationally attentive’ in the sense that she understands the incentives facing her and responds correctly in a qualitative sense, but still sometimes makes mistakes in deciding when to pay attention and when not to. Here, the specific mistake is excessively weighting the decision cost over the expected benefit of active choice. This differential weighting causes agents to fail to make an active choice when $v^* - v^d \in \left(c, \frac{c}{\beta}\right)$, situations in which a rationally-attentive agent would make an active choice. This is a mistake because the agent would receive a positive net welfare improvement from making a choice in this range. A special case of this parameterization is the beta-delta model of present-bias most commonly used in the literature on default effects in retirement saving ([Choi et al. 2003](#), [Carroll et al. 2009](#), [DellaVigna 2018](#)), if we think of agents as discounting their *future* value of insurance coverage relative to their *current* cost of making an active choice. The fact that mistakes are made on an intermediate interval is consistent with the logic in these models that agents procrastinate when the stakes are sufficiently low, but do not do so when stakes are very high. This specific parameterization does not reflect the various ways in which agents can be boundedly-rationally inattentive, but common alternatives have similar features, typically involving the agent shading or being unaware of some part of the benefit of their choices ([Ho et al. 2017](#), [Heiss et al. 2016](#)).

Random/Exogenous Attention. Finally, in brown, we plot the active choice probability for an agent for whom c is once again a constant, but for whom $\beta = 0$ and $k \sim U[c - (1 - \theta), c + \theta]$ for $\theta \in (0, 1)$. That is, $(100 \times \theta)\%$ of the time, the agent receives a transitory positive attention shock that is big enough to induce them to overcome their decision costs. For such agents, attention is entirely driven by transitory shocks, and not by any fundamentals in the value of their available options. We think of these types of agents as being irrationally or exogenously inattentive in the sense that their active decision-making is completely orthogonal to the benefit of making an active choice. The literature on inattention and salience in taxation ([Chetty et al. 2009](#)) has implicitly relied on models similar to this parameterization, in that they model agents

as having some fixed probability of, e.g. ignoring sales taxes.¹⁵

The decision rules embedded in these models are simple, but they highlight the similarities and differences between these views of active choice. While there are important distinctions between rational and boundedly-rational models in terms of welfare considerations, Figure 6a shows that typical boundedly-rational models, in terms of unequal weights on decision criteria, still are very ‘rational’ in the sense that paying attention is still increasing in the benefit of doing so.

These are, however, stylized extreme cases of potential parameterizations. One particular parameterization we have made is that cognitive costs c are fixed at some level c for *all* agents. If, instead, c is heterogeneous, we will not expect to see a ‘jump’ in aggregate active choice propensity the way we do in Figure 6a. For any given individual (or subsample of individuals) with cost c , this jump will occur, but it may vanish in the aggregate as we average over the distribution of c within a given $v^* - v^d$ cell. In Appendix Figure A8, we relax the assumption of a constant c and instead assume c is drawn from a uniform distribution $c \sim U[\bar{c} - \epsilon, \bar{c} + \epsilon]$. We plot the relationship between the probability of active choice and the $v^* - v^d$ gap for different values of ϵ , with higher values of ϵ (as illustrated by the shift from blue to red) implying more heterogeneity in c . One important note to consider is that, as the population variance in c increases, the cross-sectional relationship between stakes and active choice propensity (within a fixed range of stakes) begins to flatten, although it does still retain a small upward-sloping ‘rational’ component. We explore the implications of this insight in Section 7.

4.2 Welfare

We want to use our general framework to determine how optimal policy is affected by the features of each model. To do so, we first need to define welfare. Understanding how welfare differs across models of behavioral agents is complex.¹⁶ Thinking about welfare when default effects matter requires us to specify the potential welfare losses from the misallocation of agents into the default option, away from their most-preferred option. We measure this as $v^* - v^d$. Agents who fail to make an active choice are misallocated to the default and generate this welfare loss, while agents who make an active choice suffer no such loss. This setup mimics how structural empirical work in this literature has thought about measuring welfare (Handel 2013, Polyakova 2016).

Using our general framework from above, for any given model of active choice we can express welfare as:

$$W(v^*, v^d, c) = 1\{A(v^*, v^d, c) \geq 0\}(v^* - c) + 1\{A(v^*, v^d, c) < 0\}v^d \quad (2)$$

¹⁵As Morrison and Taubinsky (2020) point out, one can think of the parameters estimated in this way as reduced-form representations of an underlying model that looks more like a rational inattention model.

¹⁶This section is meant to highlight some topics of interest rather than serve as an exhaustive theory of welfare economics or optimal default policy under different models of choice passivity. Interested readers should turn to Bernheim and Rangel (2009) or Bernheim and Taubinsky (2018) for a more detailed treatment of how one might approach quantifying welfare when agents have potentially time- or frame-inconsistent preferences. They note that defining welfare with such behavioral agents requires normative assumptions about which time or frame is welfare-relevant. Since we do not construct any direct empirical measure of welfare in this paper, we bypass these concerns and for this section assume that there is a consistent welfare measure that has already been chosen.

When the agent makes an active decision, they pay c to receive their best alternative, whereas when they do not they receive the value of the default but pay no decision cost. Therefore, the expected welfare loss (relative to the first-best benchmark, v^*) can be written as

$$L(v^*, v^d, c) = \Pr\{A(v^*, v^d, c) \geq 0\}c + \Pr\{A(v^*, v^d, c) < 0\}E[v^* - v^d | A(v^*, v^d, c) < 0] \quad (3)$$

To build intuition, in Figure 6b we plot the *expected* welfare loss L for an agent as a function of the *potential* loss from being misallocated, i.e. the gap in value between the best plan option and the default, $v^* - v^d$. We do so for agents following the three active decision-making rules illustrated in Figure 6a.

The dark blue line once again represents the rational inattention case, with $\beta = 1$ and $k = 0$. As shown in 6a, this agent does not make an active choice when the potential welfare loss is less than c , $v^* - v^d < c$. In this range, the agent’s expected welfare loss is equal to the gap between v^* and v^d . When the gap between v^* and v^d exceeds the cost of making an active choice c , however, this agent rationally makes an active choice and receives v^* instead of v^d . In this range, the agent’s expected welfare loss is equal to the cost of making an active choice c . Importantly, the expected welfare loss is equal to c no matter how big the potential welfare loss gets because once the gap between v^* and v^d exceeds c and triggers the agent to make an active choice, the value of the default no longer matters.

The green line represents the ‘boundedly-rational’ case, with $k = 0$ as before but now with $0 < \beta < 1$. As in Figure 6a, this agent mimics the fully rational agent when the potential welfare loss is less than c , $v^* - v^d < c$. However, this agent continues to neglect to make an active choice beyond that range. Instead, this agent continues to be passive until $v^* - v^d = \frac{c}{\beta}$. This occurs because $0 < \beta < 1$ causes this agent to over-value the cost c of making a choice today relative to the benefit of that choice tomorrow $v^* - v^d$. This agent thus *mistakenly* fails to make an active choice when $v^* - v^d \in (c, \frac{c}{\beta})$. Figure 6b illustrates why this is a mistake: The boundedly-rational agent experiences a greater welfare loss than the rational agent would have if faced with the same decision. When the potential welfare loss exceeds $\frac{c}{\beta}$, however, this agent makes active choices and receives payoff $v^* - c$, as the rational agent does. This figure illustrates the idea highlighted in Carroll et al. (2009) — such an agent may be worse off with a relatively more generous default (such that $v^d \in (v^* - \frac{c}{\beta}, v^* - c)$), than with a harsher default (such that $v^d < v^* - \frac{c}{\beta}$).

The brown line represents the ‘random attention’ case, where $\beta = 1$ but each agent draws k from a distribution $U[c - (1 - \theta), c + \theta]$. Here each agent pays attention with probability $p = \theta$. Under this model, the expected welfare loss is just equal to $L = \theta c + (1 - \theta)(v^* - v^d)$. For gaps between v^* and v^d of *any* size, a random $(100 \times \theta)\%$ of agents make active choices, with those $(100 \times \theta)\%$ having a welfare loss of c and the remaining $(100 \times (1 - \theta))\%$ of agents having a welfare loss of $v^* - v^d$.

An important point illustrated by Figure 6b is that, under both the rational and boundedly-rational models, the welfare losses from agent misallocation have a strict upper bound. In the rational model, this loss is bounded by the cognitive cost c . This is a direct consequence of the ‘rationality’ at play: A rational agent will never allow themselves to suffer from a sufficiently poor state of the world that can be fixed with (costly) action. Models that allow for bounded rationality in the active choice decision relax these welfare bounds, though a bound still exists—in the present-bias model, the bound is instead $\frac{c}{\beta}$. Again, boundedly-rational agents will not allow themselves to suffer too much, although in this case they will downweight the conse-

quences of poor plan fit relative to the costliness of reparative action. In contrast, the random attention model *does not have such a bound*. Agents’ expected loss is $\theta c + (1 - \theta)(v^* - v^d)$, which is strictly monotonic in $v^* - v^d$. This arises from the fact that this agent does not respond to the state of the world at all and thus in many circumstances will simply take what comes. This important difference between the models of active choice will play a critical role in determining how optimal defaults differ across models.

4.3 Optimal Default Policy

With welfare defined and intuition in place, we now derive the optimal default under the framework we have set out.¹⁷ In general, optimal default policies will hold v^* and c fixed and change v^d in some way. In prior work, derived optimal default policy has taken on one of two forms: Either the market designer should set the default to minimize opt-outs, by trying to raise v^d as high as possible (Thaler and Sunstein 2003); or the market designer should set the default to maximize active choice, such that agents always get their most preferred option (Carroll et al. 2009). This latter design is achieved by setting v^d as *low* as possible, perhaps even punishingly low. The model we have set out allows for either of these designs to be optimal under different parameterizations. To see this we can differentiate our welfare measure with respect to v^d :

$$\frac{\partial W}{\partial v^d} = \underbrace{1 - a(v^*, v^d, c)}_{\text{Benefit to inframarginal passive agents}} + \underbrace{\frac{\partial a(v^*, v^d, c)}{\partial v^d}}_{\text{Marginal agents induced to not make an active choice}} \underbrace{E[v^* - v^d - c | A(v^*, v^d, c) = 0]}_{\text{Value of choice for marginal agents}} \quad (4)$$

Increasing the value of the default has two competing effects: First, for inframarginally passive agents, who make up a portion $1 - a(v^*, v^d, c)$ of consumers and receive the default, welfare increases by the exact amount of the increase in the value of the default. Second, some agents may respond to the increase in the default value by changing whether or not they are active choosers. Agents who become less likely to make an active choice when the default value is increased are induced to switch from making an active choice to passively following the default. Welfare for these marginal agents decreases by the difference between v^* and the value of the default v^d , minus the forgone cost of making an active choice c . In our parameterization from Section 4.1, the share of these agents in the population is given by $\frac{\partial a(v^*, v^d, c)}{\partial v^d} = \beta \cdot f(A(\cdot) = 0)$, where $f(A(\cdot) = 0)$ is the density of agents who are at the margin ($A(\cdot) = 0$) given the joint probability density of v^*, v^d, c , and k , and β reflects the relative weight put on v^d .¹⁸

Figure 7 illustrates these effects of a decrease of $\Delta v^d = v^d - v^{d'}$ in the value of the default for the boundedly rational case. As in Figure 6b, the potential welfare loss under the initial default ($v^* - v^d$) is on the x-axis and the expected welfare loss ($L(v^*, v^d, c)$ from Equation 3) is on the y-axis. The green line represents the case of the initial default, and the pink line represents the case of the new, lower-value default. It is straightforward to see that the new, lower-value default results in a welfare loss for the set of agents who fail to make an active choice under either the initial default or the new default, those for whom

¹⁷In this section, we consider optimal defaults in a partial equilibrium setting in which plan offerings do not change in response to default rules. In practice, this is unrealistic. General equilibrium considerations are likely important but beyond the scope of this paper.

¹⁸Our approach is quite similar to the approach undertaken in the optimal taxation literature. In the language of Saez (2001), our ‘inframarginal effect’ is a *mechanical effect*, while the ‘marginal effect’ is a *behavioral effect*.

$v^* - v^d < \frac{c}{\beta} + \Delta v^d$. These agents follow the default in both cases, so the lower defaults obviously lowers their welfare, with this welfare loss illustrated by the area shaded in light pink. At the same time, there is a welfare gain for agents for whom the lower default pushes their expected welfare loss high enough to overcome even their mistakenly low valuation of that future welfare loss and induces them to make an active choice. These agents have potential welfare loss in the range $\frac{c}{\beta} + \Delta v^d < v^* - v^d < \frac{c}{\beta}$. The welfare gain they experience due to the decrease in the value of the default is illustrated by the area shaded in light blue. When assessing whether a lower-value default harms or improves overall social welfare, the size of the light pink area must be compared to the size of the light blue area. Depending on the primitives, either a lower-value default or a higher-value default could be optimal.¹⁹

The intuition for the result that a higher-value default could be socially harmful is that if the policymaker is able to set a default with sufficiently low value, she can ‘shock’ all agents into making an active choice and receive $v^* - c$ in welfare. But if the policymaker does not go far enough, some inframarginal passive agents will suffer due to the decrease in the value of the default plan, in which they will ultimately enroll. In such a trade-off environment, the planner must generally decide between two extremes: A ‘benevolent’ default, which provides the maximum achievable default value, or a ‘shocking’ default, which provides $v^* - c$, although this latter piece assumes that a default sufficiently low to induce universal choice is feasible.²⁰

Note, however, that the existence of this trade-off, specifically the existence of the negative (second) effect, depends critically on two conditions. First, it requires $E[v^* - v^d - c | A(v^*, v^d, c) = 0] > 0$, which is equivalent to the idea that, on the margin, agents who fail to make an active choice are experiencing a welfare loss. While this seems like a trivial point, we note that it is *not* true for rationally-inattentive agents. For an agent with the ‘rational inattention’ parameterization $A(v^*, v^d, c) = v^* - v^d - c$, note that

$$E[v^* - v^d - c | A(v^*, v^d, c) = 0] = E[v^* - v^d - c | v^* - v^d - c = 0] = 0 \quad (5)$$

and therefore there is no trade-off in the optimal default policy and thus no benefit to a shocking default. More boundedly-rational models of decision costs do, however, induce ‘mistakes at the margin.’ In Figure 6b, the vertical red dashed line at $v^* - v^d = \frac{c}{\beta}$ illustrates the mistake at the margin, which is equal to $E[v^* - v^d - c | \beta(v^* - v^d) - c = 0] = \frac{1-\beta}{\beta}c > 0$. This is true for *any* value of v^*, v^d .

The existence of these mistakes at the margin is a necessary condition for the negative effect in Equation (4) to be non-zero, but it is not sufficient. The second necessary condition is that $\frac{\partial a(v^*, v^d, c)}{\partial v^d} < 0$, i.e. that increases in the value of the default option induce some agents to become passive. In other words, if there are no agents on the margin of making an active choice or not, it does not matter for policy whether the (non-existent) marginal agents are making mistakes. While it is likely (though not guaranteed) that there will be marginal agents when attention is rational, it need not be true more generally. When $\beta = 0$ in our

¹⁹We refer to a lower- or higher-value default as “optimal” here, but we would only know for sure that it is optimal in a very local sense. It may be that a slightly higher-value default improves welfare, but the globally optimal default is actually a lower-value one. When we take the insights of the model to the data in Section 6, we assess all policy-relevant defaults (all plans in the choice set) to address this issue.

²⁰The planner’s ability to set default values to be very high or very low is likely to be constrained. Information asymmetries (i.e., not knowing the preferences of agents) restrict a planner’s ability to match agents to their optimal plans. Political constraints may limit the planner’s ability to set harsh defaults—for example, taxing passive agents is unlikely to be an acceptable policy. The ‘shocking’ policy studied in Carroll et al. (2009), for example, merely changed the *framing* of the consequences of a failure to make an active choice rather than imposing any additional negative consequences.

parameterization above (e.g. in the ‘random attention’ case), $\frac{\partial a(v^*, v^d, c)}{\partial v^d} = 0$ *always* and thus there is no trade-off.

Generating optimal default policy with no trade-off is clear: When there is no harm to marginal agents, improved defaults benefit inframarginal passive agents at no cost and thus the optimal default should be to minimize opt-outs by maximizing v^d . On the other hand, when there is a trade-off, it may be optimal to ratchet down the value of the default.²¹ Prior work, when evaluating this trade-off, has focused primarily on considering the extent to which agents are making a mistake (on the margin and otherwise) when they do not make an active choice. This is an inherently *normative* question. [Bernheim et al. \(2015\)](#) and [Goldin and Reck \(2020\)](#) show that different normative assumptions about whether choice passivity is a ‘mistake’ or a response to real decision costs can generate both of the two extreme optimal default policies we describe above. While having a specific model of behavior is important, our exercise in this section suggests that simply making an assumption about the extent of the mistake is sufficient for optimal policy even without a fully-specified model.

More importantly, this prior work has restricted itself to model domains in which $\frac{\partial a(v^*, v^d, c)}{\partial v^d} < 0$ by assumption.²² If this is not true, and agents do not adjust their passivity (or make only very minor adjustments) in response to differences in the value of the default option, *then the question of whether or not their passivity is a ‘mistake’ is moot*. We have little prior evidence to answer that question: Are there any marginal beneficiaries in reasonable ranges of v^d ? If one sets a shocking default, how much (if any) active choice does it actually induce? In the next sections we attempt to provide evidence on this question by asking whether beneficiaries in our setting who face relatively worse defaults are indeed more likely to opt out.

5 Should Beneficiaries Care About Default Rules?

As we described in the previous section, optimal default policy depends on the extent to which beneficiaries respond to changes in the value of the default by making (or not making) an active choice. In Sections 2 and 3 we showed that the vast majority of LIS beneficiaries follow defaults. However, we have not yet explored differences in the value of the various defaults offered to LIS beneficiaries. Further, we have not studied the relationship between the value of these defaults and the tendency to make an active choice, $\frac{\partial a}{\partial v^d}$, the critical parameter for determining the optimal default according to the model presented in Section 4.

In this section we attempt to establish differences in the value of the various defaults. We show that the defaults differ significantly in the consequences for LIS beneficiaries.²³ Specifically, we show that some

²¹Additional considerations include the extent to which the policymaker can achieve a v^d close to v^* by correctly picking a beneficiary’s welfare-maximizing plan option; if they can get very close, a ‘paternalistic’ default may be optimal even in the presence of a trade-off, because at such high values of v^d , if $c > 0$, then $E[v^* - v^d - c | A(v^*, v^d, c) = 0] < 0$. When this is not true, and the policymaker would prefer to ratchet down the value of the default, extreme ‘shocking’ defaults can arise as an optimum because the share of inframarginal agents who are hurt by such a default decreases as the default becomes more harmful and agents become ‘choosers,’ and thus $\frac{\partial W}{\partial v^d}$ may *grow* in magnitude as v^d falls.

²²[Goldin and Reck \(2020\)](#) specifically employ a model akin to our ‘rational inattention’ model with $A(\cdot) = v^* - v^d - c$, but allow c to have welfare-relevant and welfare-irrelevant components. The size of the welfare-irrelevant component of c in their model is equivalent to the size of what we call the ‘mistake on the margin.’

²³A potential explanation of our results from prior sections could be that, in a rational attention model, if the choices we study are low-stakes in that plans differ only in minute ways, we would expect high default propensities. This is rejected by our results in

defaults result in large losses of valued drug consumption while other defaults leads to smaller consumption losses. We then leverage random assignment of default plans to continuing LIS beneficiaries whose plans lose benchmark status to estimate $\frac{\partial a}{\partial v^d}$ and to assess whether $\frac{\partial a}{\partial v^d}$ is small enough to suggest that the optimal default is a paternalistic one that attempts to nudge beneficiaries to plans that match their drug needs or large enough to suggest that a shocking default that attempts to trigger active choices might be optimal.

We study differences in the value of the various defaults at play in the LIS program in two steps. First, we study differences in the value of the default of remaining in one's incumbent plan versus reassignment to a randomly selected default benchmark plan. Second, we study differences in the value of each of the plans in the set of plans from which the new default plan is randomly chosen.

It is not immediately obvious that LIS beneficiaries would value the various defaults differently from one another. In other words, it is not obvious that LIS beneficiaries truly care which benchmark plan they are enrolled in. In prior work ([Abaluck and Gruber 2011](#), [Handel 2013](#), [Polyakova 2016](#)), the cost of following the default is enrolling in a plan whose premium and cost-sharing parameters are poorly aligned with the enrollee's expected drug or healthcare consumption, possibly resulting in worse financial outcomes for beneficiaries. In contrast, for the LIS beneficiaries we study, premiums and cost-sharing amounts are highly subsidized and uniform across plans, conditional on the inclusion of a drug on the plan's formulary. With no differences in premiums and cost-sharing, it may seem like beneficiaries might not care at all which plan they are enrolled in, and it would thus not be surprising that they have such a strong tendency to follow defaults.

However, a lack of differences in cost-sharing and premiums does not imply that there is no difference in the beneficiary's experience. Formularies still determine what drugs are and are not covered by a plan. As we show in [Appendix B](#), benchmark plans do vary significantly in what portion of a beneficiary's drugs they cover. Such differences suggest that beneficiaries are likely to receive different levels of utility from different default options. But how can we measure those differences in utility?

Naturally, directly measuring utility under each default option is not possible in this setting. Our approach instead asks whether beneficiaries face *different consequences* of following each default. Specifically, we ask whether beneficiaries are able to maintain the same level of drug consumption under the two defaults. As long as beneficiaries value the consumption they get from their Part D plan—highly likely, given that our sample is composed of low-income, vulnerable individuals, many of whom are in poor health—changes in consumption will be informative about beneficiaries' welfare loss from following the default. We think of this approach as an ex-post effect of following the default, in the spirit of other recent work measuring welfare effects of health insurance policies based on their effects on consumption and health spending ([Finkelstein et al. 2019](#)). This approach stands in contrast to the ex-ante approach of prior work in health insurance, where researchers determine the quality of a match between a consumer and their plan according to a parametric model of health insurance demand. We use this approach both for (1) assessing differences in the value of the default of remaining in one's prior plan versus reassignment to a randomly-selected default plan and (2) assessing differences in the value of the various default plans to which a beneficiary could be randomly assigned.

this section.

5.1 Value of Remaining in Prior Plan Versus Reassignment

To assess differences in the consequences to beneficiaries of following the default of remaining in their incumbent plan versus following the default of re-assignment to a randomly-selected default plan, we again leverage the natural experiment from Section 3. Recall that when a year $t - 1$ benchmark plan submits a year t premium bid exceeding the (ex-ante unknown) year t subsidy, the default for all beneficiaries enrolled in that plan in $t - 1$ changes from remaining in the plan to reassignment to a randomly-selected default benchmark plan. Beneficiaries in a $t - 1$ benchmark plan whose year t bid is below the year t subsidy maintain the default of remaining in their year $t - 1$ plan. We use a stacked difference-in-differences design to compare a variety of outcomes pre- vs. post-assignment for beneficiaries in plans that bid above the subsidy (treatment group) versus beneficiaries in plans that bid below the subsidy (control group).²⁴

For each market (Part D region) and year we identify all beneficiaries who were enrolled in a benchmark plan in year $t - 1$ that submitted a year t premium bid within \$6 of the year t subsidy, in the spirit of the RD design. We construct a dataset including all of these beneficiaries, with beneficiaries in plans bidding above the subsidy being classified as “treatment” and beneficiaries in plans bidding below the subsidy being classified as “control”. We assign an “experiment ID” d to each market-year dataset, as each dataset includes all of the necessary ingredients to run a single difference-in-difference regression. We then stack all of these datasets and run the following regression pooling across all of the individual difference-in-difference experiments:

$$y_{itd} = \beta(\text{BenchmarkLoss}_{id} \times \text{Post}_{td}) + \gamma_{id} + \eta_{td} + \epsilon_{itd} \quad (6)$$

where y_{itd} is the outcome of interest for person i in time t in experiment d , $\text{BenchmarkLoss}_{id}$ is a dummy variable equal to one if person i is in the treatment group in experiment d , i.e. that their incumbent plan submitted a bid which would cause it to lose benchmark status, and Post_{td} is a dummy variable equal to one if quarter t is after the change in default occurs in experiment d . γ_{id} is an individual-by-experiment fixed effect, recognizing that an individual may appear in multiple experiments. η_{td} is a full set of (event-)time-by-experiment fixed effects, and ϵ_{itd} is a random error term. The inclusion of the time-by-experiment fixed effects η_{td} ensures that this regression specification compares changes in outcomes for beneficiaries who are treated (i.e. re-assigned) in a given year only to changes in outcomes for beneficiaries who were enrolled in plans that submitted bids within \$6 of the benchmark subsidy (and were thus *almost* treated) in that same year. This accomplishes two goals. First, it ensures that the treatment and control beneficiaries are as comparable as possible. Second, it eliminates the problem of the more typical two-way fixed effect model pointed out by [Goodman-Bacon \(2019\)](#), [Abraham and Sun \(2020\)](#), and [de Chaisemartin and D’Haultfoeuille \(2020\)](#) where dynamic treatment effects (which we might expect here) may lead to biased estimates of

²⁴We use the difference-in-differences design here rather than the RD design used in Section 3 in order to maximize power for our statistical tests of the consequences of the change in the default. The outcomes we are studying here (drug consumption) are noisy, requiring much more statistical power than the less-noisy enrollment outcomes studied in Section 3. The assumptions necessary for the difference-in-differences design to be valid are stronger than the RD assumptions, but we provide convincing evidence below that these assumptions are satisfied here. Moreover, we can see in Figure A5 that the relationship between active choice and the running variable (incumbent plan premium bid relative to the benchmark) is close to flat, so we are not worried about differences in active choice propensity near the threshold versus farther away, especially once we absorb beneficiary fixed effects.

treatment effects due to 2x2 comparisons between “late-treated” and “early-treated” units. Here, we only compare currently treated beneficiaries to beneficiaries who (1) were in plans that bid just below the subsidy in the same year that the treated beneficiaries’ plans bid just above the subsidy, and (2) were untreated throughout the contemporaneous 8 quarters pre-treatment and 4 quarters post-treatment, eliminating the late vs. early comparison. One potential quirk in our approach is the fact that beneficiary-years can appear multiple times in the dataset on which we estimate our regression. To deal with this, we cluster at the level of the beneficiary, rather than the beneficiary-by-experiment level. Our final sample is composed of 464,557 person-by-experiment observations, over 12 quarters.

Finally, we note that while the regression from Equation 6 will estimate the effect of the default of random re-assignment versus the default of remaining in one’s incumbent plan, it will not estimate the effect of switching from one’s incumbent plan to a randomly chosen plan. This is because beneficiaries are under no obligation to follow the default. In Section 3.1, however, we showed that the vast majority (over 90%) of treatment group beneficiaries do indeed follow that default and switch. This suggests that our ‘intent-to-treat’ estimates are likely quite close to the average treatment effect.

Results. Our primary outcome measure, y_{itd} , is the log of the beneficiary’s total covered drug spending in a quarter. This spending includes both what the beneficiary paid out of pocket (typically very little) as well as what Medicare and the insurer paid on their behalf. Panel (a) of Figure 8 shows the event study version of Equation 6 for overall drug spending.²⁵ We can see that the change in the default from remaining in one’s incumbent plan to forced assignment to a different, randomly-selected default benchmark plan does indeed have a measurable effect on drug spending, causing it to suddenly plummet for beneficiaries enrolled in plans that lost benchmark status starting in quarter T , the first quarter after the change in the default. Encouragingly, we do not see any differential changes over time between beneficiaries in plans that lost benchmark status versus those that did not in years *before* the benchmark status loss, providing evidence that the parallel trends assumption is satisfied here. In the first column of Table 3, Panel A, we present our estimate of β from Equation 6, where we pool across all quarters in the post-period. There is a clear negative effect of the new default on drug consumption, with the loss in benchmark status and the subsequent change in the default resulting in a 6.4% drop in drug spending for affected beneficiaries, approximately \$213 off a base of \$3,329. This represents a nontrivial loss in valued consumption. In Appendix C we show that this reduction persists at least 24 months after the implementation of the new default, implying that the mismatch is permanent in nature and not addressed by a beneficiary “adapting” her drug use to the formulary of her new default plan.

One concern is that this analysis includes both changes in the drugs consumed and the price of the drug consumed. Different Part D plans are able to negotiate different prices for identical drugs (potentially leading to their different premium bids), and we do not want to confound this with consumption, especially given the fact that beneficiaries do not pay for the drugs themselves. Therefore, we replicate our estimation of β for an alternative measure of drug spending, where we replace plan-specific unit prices for a given drug (identified at the NDC level) with that drug’s average price across all plans in our sample for that year. The

²⁵The difference between this and the specification in Equation 6 is that, rather than have a single β for distinguishing between pre-treatment and post-treatment periods, we estimate a treatment effect for every period, including pre-treatment periods.

full event study for this price-normalized outcome is presented in Panel (b) of Figure 8. The estimate of β from Equation 6 is found in the second column of Table 3, Panel A. While the effects are slightly smaller (5% reduction vs. 6.4%), they remain qualitatively similar and important.

Another concern is that our spending results also partially reflect induced substitution between more- and less-expensive drugs meant to treat the same ailment. The price of a drug may genuinely reflect its value to the patients, but it also may simply reflect differential market power. We construct an additional normalization to handle this case: We normalize prices at the drug *therapeutic class* level. This is an extreme normalization, as drugs within the same therapeutic class may be poor substitutes for one another. In column 2 of Table 6, Panel A, we see that, with this normalization, effects are still present, but smaller, with the default change resulting in a 2.1% drop in normalized spending. We think of these two normalization approaches as representing a bounding exercise between two extremes: On the one hand, taking our drug-level price normalization seriously for welfare purposes requires the assumption that drug prices are perfectly representative of drug value; on the other hand, our class-level normalization approach assumes that all drugs within the same therapeutic class are perfect substitutes and all within-class price differences are arbitrary.²⁶ As far as we know, there are no general elasticity of substitution estimates for drugs that would inform this question.²⁷

Another way to evaluate the value of the foregone drugs is by looking at the effects of the new default on consumption of specific categories of drugs that physicians deem to clearly have high value for patients. In this spirit, we characterize drugs as high-value in a clinical sense, based on whether they belong to therapeutic classes targeting key chronic conditions; these conditions include diabetes, depression, hypertension, and asthma. The regression estimates for this outcome are found in Table 4. We find that the reduction in spending on high-value drugs is almost identical to the overall reduction in drug spending. Thus, lost drugs are not simply less-beneficial ones. We also estimate an effect for drugs that treat chronic illnesses. The disruption of consumption of these drugs in the process of the transition from the incumbent plan to the randomly-selected default may be more costly for the beneficiary. Such a disruption would be even more likely if the drug faces new authorization restrictions under their new plan. We define this category as drugs for which the median number of annual, per person fills is greater than two, among the subset of people with at least one annual fill (Einav et al. 2015). Drugs treating chronic illnesses see much larger decreases (7.8%) than other drugs (2.8%).²⁸

Finally, we estimate effects on non-drug medical consumption, a potential proxy for beneficiary health (Chandra et al. 2010). Medical spending is covered by Medicare Parts A and B, which are distinct from Part D, so Part D plan assignment does not affect non-drug spending mechanically. The third column of Table A2 details the output from this regression. While we find that reassignment increases non-drug spending,

²⁶There is also the potential for match-specific quality, where a given beneficiary has value for a specific drug rather than another. This private value will not be fully captured in the price and would make our spending effect results an underestimate of the true welfare effects.

²⁷In Table A2 we also measure the effects of the default change on the total number of prescription fills and days supply consumed. These effects are relatively small; however, interpreting them requires the extreme assumption that a unit of *any* drug has equivalent value. Given that the managed care restrictions used by plans are typically imposed specifically on high-cost drugs, this assumption seems unrealistic.

²⁸One important point to note is that these results cannot be directly compared to other results on how drug spending elasticities differ across different types of drugs, such as in Brot-Goldberg et al. (2017). In our setting the estimates include both the behavioral response *and* instrumental differences in how different plans cover and restrict different drugs.

the effects are small and noisily estimated. Further, this estimate may not fully capture the effect on non-drug spending given that this effect could be expected to arrive with a significant lag and manifest over the longer-term.

Our broad takeaway from this analysis is that beneficiaries face real consequences from following the default of being randomly auto-assigned to a new benchmark plan rather than remaining enrolled in their incumbent plan. They are thus unlikely to be following this default out of indifference, as they are forced to reduce their consumption of high-value drugs.

5.2 Differences in Valuation of Randomly-Selected Default Plans

We now ask whether there are different consequences of being assigned one default benchmark plan versus another. Specifically, we assess whether beneficiaries face different consequences of being assigned to a plan that covers more versus fewer of their drugs. We measure ‘fit’ of a plan’s formulary as the share of the prescriptions taken by the beneficiary prior to the change in the default that are covered by the plan. Plans that fail to cover a beneficiary’s drugs may result in hassle costs of switching, as well as reduced welfare and consumption. For each beneficiary who is reassigned, we compute her fit for each benchmark plan in her market in her year of reassignment. An absolute measure of fit is a function of the number of drugs taken and thus *not* necessarily exogenously assigned. Instead, for each beneficiary, we rank plans in the beneficiary’s choice set and use the *within-person* quintile of fit for the beneficiary’s assigned default plan, relative to other benchmark plans. The probability of being assigned a plan in any given quintile is $\frac{1}{5}$, which is independent of the beneficiary’s characteristics.

To estimate differences in the consequences of being assigned a poorly-fitting default versus a better-fitting default, we replicate our stacked difference-in-differences design, but interact the indicator for a beneficiary’s incumbent plan losing benchmark status with the fit quintile of their randomly-assigned default plan:

$$y_{itd} = \delta_t(WorstFit_{id} \times BenchmarkLoss_{id}) + \beta_t BenchmarkLoss_{id} + \gamma_{id} + \eta_{td} + \epsilon_{itd} \quad (7)$$

In the regression specification, δ_t represents the differential effect of being assigned to a plan in the worst-fitting quintile on drug consumption outcomes and β_t represents the impact of being reassigned to a plan generally. Random assignment of the fit quintile allows us to identify δ . Figure 9a graphs the event-study estimates of $\beta_t + \delta_t$ (the total effect of being assigned a default plan in the worst-fitting quintile) and β_t (the effect of being assigned a default plan in any other quintile), with the outcome as the log total drug spending for a beneficiary in a given quarter. Beneficiaries assigned high- versus low-fitting default plans had similar levels of drug spending prior to re-assignment, as expected given random assignment of default plans. After the new default is implemented, however, there is a much larger spending reduction for those assigned low-fitting defaults relative to the reduction experienced by those assigned high-fitting defaults. Figure 9b presents the event study where drug prices are normalized across plans, revealing that significant differences remain between the consequences of assignment to a low-fitting versus a high-fitting default plan. Panel B of Table 3 presents this regression run at the yearly level, including log spending (column 1) as well as log price-standardized spending (column 2). Re-assignment to the worst-fitting plans reduces

drug spending by 12.6%, compared to a 4.3% effect of being reassigned to any other plan. Statistically and economically significant differences remain as our various price normalizations are applied (see columns 2-4).²⁹ In Appendix C we show that these differences in consumption persist for at least *two years*, implying that beneficiaries are not adapting to their new formularies over time.

Thus, in all cases we see real effects on quantities from poor formulary fit, implying that beneficiaries are forgoing consumption rather than merely switching prescriptions. Under the assumption that beneficiaries value their drug consumption, which is reasonable for low-income individuals in relatively poor health, this provides evidence that beneficiaries face different consequences from each of these default options and are thus likely to value these defaults differently.

6 Do Worse Defaults Drive Beneficiaries to Make Active Choices?

In Section 4, we showed that the optimal default policy depends on the extent to which beneficiaries respond to changes in the value of the default by making (or not making) an active choice, $\frac{\partial a}{\partial v^d}$. In Sections 2 and 3 we provided convincing evidence that LIS beneficiaries follow defaults. We have now also established that there are real negative consequences to beneficiaries of following some defaults, in the form of large losses of drug consumption. Despite these negative consequences, an overwhelming portion of beneficiaries follow those defaults. This begs the question of what these results imply about $\frac{\partial a}{\partial v^d}$ and the optimal default in this setting.

We first point out that, when taken together, two of our previous results suggest that $\frac{\partial a}{\partial v^d}$ is likely to be quite small: (1) when beneficiaries' default is changed from remaining in their prior plan to re-assignment to a randomly-selected default benchmark plan beneficiary drug consumption decreases between 2 and 6% on average; and (2) over 90% of beneficiaries follow that default, and less than 5% of beneficiaries are induced by the imposition of that default to make an active choice. This suggests that any potential detrimental effect of paternalistic defaults on marginal choosers is also likely to be small. However, we can go further and directly estimate $\frac{\partial a}{\partial v^d}$ by leveraging the randomly-assigned defaults.

To do so, we attempt to assess whether *differences* in the consequences of not making a decision are related to *differences* in active choice propensity. Doing this is nontrivial. First, we need to find variation in v^d that exists *independently* from differences in v^* and c , which might otherwise affect active choice propensity. Second, v^d is a function of beneficiaries' preferences and is thus inherently unobservable, so we cannot directly measure active choice as a function of it.

To generate exogenous variation in v^d , we leverage the fact that in the setting we describe in Section 3 defaults are *randomly assigned* to beneficiaries whose year $t - 1$ benchmark plans bid above the year t subsidy. When beneficiaries' plans lose benchmark status, they are assigned a *randomly-selected* default

²⁹If beneficiaries assigned low-fitting default plans are more likely to opt out of that default relative to beneficiaries assigned to high-fitting plans, the differences in drug spending reductions could be due to different plan-effects on spending or to differences in the population of 'compliers' in the low- and high-fit cases. Two arguments suggest that this is not an issue here and that these estimates reflect only differences in plan effects. First, we show in Section 6 that differences in opting out of high- versus low-value defaults are fairly trivial. Second, we would expect any differences in opting out to push against finding a larger spending reduction from the low-fitting defaults. We would expect that the marginal beneficiaries induced to opt out of the lower-fitting default would have larger potential treatment effects (thus, why they opt out) relative to inframarginal beneficiaries. The absence of these beneficiaries from the complier population for the low-fitting plans would thus result in lower estimates of spending reductions.

plan. Since assignment is random, and orthogonal to beneficiary characteristics, it generates differences in v^d that are orthogonal to v^* and c . Importantly, each person’s randomly chosen default plan is revealed to them months prior to the implementation of the new default. If the plan were not revealed prior to the implementation of the new default (i.e. prior to the forced switch taking place), then v^d would be the expected valuation across all plans that the beneficiary could potentially be assigned to, and there would be no variation in v^d within a market. However, since the defaults are revealed months prior to when they are implemented, we can think of each person’s assigned plan as their own randomly chosen person-specific default and thus use the *random* variation in v^d to identify $\frac{\partial a}{\partial v^d}$.

However, even with this exogenous variation in the default itself, we still do not have a direct measure of v^d , since we have no measure of v more generally. We thus take what we call a ‘proxy approach.’ If $v_i^j = \gamma x_i^j + u_i^j$, where x includes some observable characteristics of the plan j (potentially modified by the characteristics of the beneficiary i) and u contains all other joint factors, then $\frac{\partial a}{\partial x^d} = \frac{1}{\gamma} \frac{\partial a}{\partial v^d}$. So, as long as variation in x^d provides sufficient variation in v^d (i.e., as long as γx^d is an important component of how beneficiaries value d), and γx^d has non-negative correlation with other features u^d , we can use variation in a ‘proxy’ characteristic to recover $\frac{\partial a}{\partial v^d}$.³⁰ The relative importance of x versus u is also important for generalizability: If x matters some but not much for valuation, then this method will only leverage trivial variation in v^d to identify $\frac{\partial a}{\partial v^d}$, raising questions about any conclusions drawn from our specific estimates.

For our proxy characteristic, we use the beneficiary-specific ranking of the fit of the assigned plan relative to the fit of the other benchmark plans in the beneficiary’s choice set (i.e. the set of plans from which the beneficiary’s randomly-selected default was chosen). As we showed in Section 5.2, this beneficiary-specific fit ranking captures real differences in the effects of the default on beneficiary drug consumption, with low-fitting plans causing a loss in drug consumption at least three times as large as the loss caused by higher-fitting plans. Under the assumption that beneficiaries value their drug consumption, which is reasonable for low-income individuals in relatively poor health, this provides evidence that our variation in x^d (fit) proxies for substantial variation in v^d .

Given that we have established that the randomly assigned defaults with different levels of fit provide real exogenous variation in v^d , we proceed to leverage that variation to estimate $\frac{\partial a}{\partial v^d}$. In Figure 10a, we plot the share of beneficiaries who switch from their incumbent plan (i.e. the share who follow the default) as a function of whether they were assigned a default plan in the worst-fitting quintile. As shown in section 3, overall only around 3% of beneficiaries opt out of the default to remain in their incumbent plan. More importantly for $\frac{\partial a}{\partial v^d}$, we also see little heterogeneity in this rate across individuals facing heterogeneous consumption losses. This can be seen more clearly in column (1) of Table 5, where we present results from the stacked difference-in-differences specification from Equation 7 but where the outcome is a dummy variable for switching out of the incumbent plan: The coefficients in Panel B of the table suggest that, while the overall effect of the change in the default from remaining in one’s incumbent plan to reassignment to a randomly-selected default plan is to increase the probability of switching by 95.5 percentage points, those assigned to low-value defaults are slightly *more* likely to switch plans, albeit only by a statistically-

³⁰We can think of x as an implicit ‘instrument’ of sorts for v . Caplin et al. (forthcoming) use a similar logic to show how one can use this variation to trace out an attention ‘supply curve’ in a rational inattention model. $\frac{\partial a}{\partial v^d}$ is the slope of that curve. The empirical approaches of Brown and Jeon (2020) and Morrison and Taubinsky (2020) are similar to ours.

insignificant 0.2 percentage points. This result suggests that $\frac{\partial a}{\partial v^d}$ is likely to be close to zero in this setting, despite real variation across defaults in the effects on an important beneficiary outcome, drug consumption.

In Figure 10b we consider the possibility that while beneficiaries assigned to worse plans are no more likely to opt out of the default and remain in their incumbent plan, these beneficiaries may be more likely to be triggered to actively choose some other non-default plan. Here, we plot the cumulative portion of beneficiaries who have actively opted out of their default in each quarter after the change in the default in quarter T . We present the output from a regression of choice status by the end of the following year on fit in column (2) of Table 5. As before, very few beneficiaries make active choices. While beneficiaries assigned to a low-fitting plan are slightly more likely to make an active choice relative to beneficiaries assigned to high-fitting plans, the difference is fairly trivial (less than 2% of beneficiaries) and in both cases fewer than 7% of beneficiaries make an active choice. This remains true all the way through the end of the plan year, ruling out the possibility that beneficiaries assigned to poor-fitting plans only come to realize it after the plan year begins and then respond accordingly. The overall active choice rate barely budges between quarter T and quarter $T + 3$. In Appendix C, in Figures A17 and A18, we show that these patterns persist in our restricted two-year sample. This is all true despite the group assigned poorly fitting default plans experiencing *three times* the consumption losses when they follow those defaults. This is a surprising non-response to a highly salient treatment: Unlike in the retirement savings context, here beneficiaries realize that a change has occurred when they arrive at the pharmacy and realize that their prescriptions are no longer covered. Yet there is almost no relationship between the probability of such an event and the beneficiary’s propensity to opt out of her default and make an active choice, either initially or even after the consumption loss has occurred. Moreover, because the beneficiaries had experience in a prior plan, they can pay a small price (a few dollars per month) to return to that plan and regain full coverage of their previous prescriptions, incurring virtually no search costs. Yet almost no beneficiaries do this, no matter how damaging the new plan is in terms of fit. Again, this suggests that $\frac{\partial a}{\partial v^d}$ is close to zero.

To aid with interpretation of these results, in Figures 11a and 11b we group beneficiaries together by the quintile of fit of their randomly-assigned default plan, and we plot their estimated one-year plan switching and opt-out propensities against their estimated drug consumption losses from remaining in their assigned plan.³¹ This is an empirical analogue to Figure 6a. We can think of the consumption reductions as an analogue to the gap between v^* and v^d , with larger consumption reductions corresponding to a larger potential welfare loss (x-axis of Figure 6a). It is clear that as the consumption effect rises, there is effectively no change in the propensity to switch out of the incumbent plan (Figure 11a) or the propensity to make an active choice (Figure 11b), consistent with the prior figures, and consistent with the random attention model of consumer choice.

One possible alternative explanation is that low benefits from choice (small differences between v^* and v^d) could rationalize these results. That we find significant differences in consumption is evidence against such a conclusion. Further, in Appendix Figures A10 and A11 we re-produce Figures 11a and 11b, restricting to beneficiary types where the “spread” in fit between the worst and best fitting plans is large. For some beneficiaries, this spread is quite small, and for these types the variation in fit we are using to identify $\frac{\partial a}{\partial v^d}$ is limited. Thus, we restrict to beneficiaries with a large spread in fit across available

³¹Quintile-specific estimates of switching, active choice, and drug spending effects are found in Appendix Table A5.

plans in order to assess whether the key results hold for this group. For each beneficiary, we measure the variance of their fit measure across the benchmark plans available to them. We then rerun our main specification for a subsample including only the top 50% and 25% of treated beneficiaries, ranked by this variance. The consumption effects of being assigned to a high- versus a low-fitting plan are even starker for these beneficiaries: Misallocation to a poorly-fitting plan decreases drug spending by 15 percentage points relative to well-fitting plans. However, there is still little relationship between these larger spending effects and the propensity to make an active choice. In all cases, fewer than 10% of beneficiaries make an active choice, even when assigned default plans that cause substantial consumption reductions.

In summary, our results in this section suggest little or no relationship between the value of the default and the propensity for a beneficiary to make an active choice. This is true even for beneficiary types for whom the variance in value of the defaults is large. Further, this evidence comes from *randomly-assigned* defaults, eliminating standard selection-related alternative explanations for our findings. Ultimately, it appears to be the case that $\frac{\partial a}{\partial v^d}$ is close to zero in our setting, implying that the benefits of paternalistic defaults that match beneficiary preferences are highly likely to outweigh any (small) cost from causing (few) beneficiaries to be less likely to make an active choice.

Beyond the clear implications for the optimal default in this setting, our results provide new insights into which of the decision cost models best explain consumer behavior. A model that explains these results must have the feature that the aggregate $\frac{\partial a}{\partial v^d}$ is close to zero. As we described in Section 4.1, a model in which attention is exogenous/random explains such a result. However, we also showed that models of rational or boundedly rational inattention with sufficient heterogeneity in decision costs c (given by the red line in Figure A8) also predict such a finding. In the final section, we try to disentangle these two explanations and provide additional insights into consumer decision-making behavior.

7 Unobserved Heterogeneity in Attentiveness

Our results in the previous section suggest that variation in the potential benefits of active choice do not seem to drive variation in the rate of active choice. This leaves a puzzle: What *does* explain this variation? Recall that at any given point in time, while two-thirds of beneficiaries are in a plan they did not choose, one-third of beneficiaries did in fact actively choose their current plan. So why do some people make active choices and others do not? Our remaining explanation is that beneficiaries vary in their level of attentiveness. In this section, we disentangle two types of variation in attentiveness. The first is that beneficiaries are of different ‘types,’ some of whom are generally more attentive (or face lower costs c of making a decision) and some of whom are less attentive (or face higher costs c of making a decision). The second is that beneficiaries receive transitory shocks to their attention from time to time, and when they receive a positive shock, they make a decision.

We note that under both of these explanations, our previous result that there is little relationship between v^d and active choice still suggests that in our setting a paternalistic default matched to beneficiary preferences is likely to be welfare-improving relative to a shocking default purposely mis-matched to beneficiary preferences. Thus, it may seem like these two explanations for this lack of a relationship are a distinction

without a difference. We point out, however, that such a distinction is important for thinking about default policy design and the distributional impacts of defaults. Under the first explanation, passivity is concentrated among a subset of beneficiaries, which might encourage targeted interventions. Further, universal paternalistic defaults would provide the most benefit to these passive types and little benefit to other types. On the other hand, under the second explanation, *all* beneficiaries are potential non-choosers, suggesting that applying paternalistic defaults to *all* beneficiaries may be helpful and that all beneficiaries would benefit to some extent from these defaults.

In this section, we decompose this variation. As a benchmark, consider a simple model of latent attention (borrowing the A notation from our framework in Section 4):

$$A_{it} = c_i + k_{it}$$

where c varies across beneficiaries i according to some distribution, and k varies i.i.d. across time t within a beneficiary. The ideal dataset to measure this would be panel data in which we measure A . Since A is latent and thus unmeasurable, we can instead proxy for it with a , an indicator for whether the beneficiary makes an active choice. But we can only observe active choice when the beneficiary faces a default of auto-assignment, which is rare in our data consisting of many months of beneficiary choices, few of which involve the default of auto-assignment (vs. continued enrollment in the beneficiary’s incumbent plan). The default of auto-assignment is especially rare for beneficiaries who are actively enrolled in their plans and thus do not face the risk of being reassigned as a result of benchmark status loss.

Instead, we use an alternative natural experiment that occurs in our data. When a plan exits at the end of a year, *all* of its enrolled beneficiaries are randomly assigned default plans from the set of benchmark plans available in the market in the following year. Recall that in the case of plans bidding above the benchmark, the switch to the randomly-chosen default plans only applies to beneficiaries previously auto-assigned to the plan. Here, when the plan exits the market entirely, *all* enrollees are allocated a randomly-chosen default plan, not just those previously auto-assigned. Thus, in the year after exit, we are able to directly measure active choice for all relevant beneficiaries.

For beneficiaries in exiting plans we therefore have two pieces of data on their active choice propensity: (1) their method of enrollment (active vs. passive) for the exiting plan and (2) their method of enrollment in their next plan post-exit. If these decisions are highly correlated, i.e. if beneficiaries who had actively enrolled in the exiting plan are likely to make an active choice post-exit, that suggests that the variation is better-explained by individual-specific heterogeneity than transitory shocks. To see this, note that $\text{Cov}(A_{it}, A_{it'}) = \text{Var}(c_i) + \text{Cov}(c_i, k_{it}) + \text{Cov}(c_i, k_{it'}) + \text{Cov}(k_{it}, k_{it'})$, which, due to the assumed independence of k across time (and the independence of c and k by construction), is equal to $\text{Var}(c_i)$. Therefore, the propensity of prior active choosers to continue to make active choices allows us to identify the magnitude of the variance component of the individual-specific component relative to the transitory component.

We construct a subsample of all beneficiaries whose plan exited within the time window covered by our data. We restrict only to beneficiaries enrolled in benchmark plans, as there cannot be beneficiaries enrolled in non-benchmark plans who did not actively enroll in those plans.³² This gives us a sample of

³²We allow our sample to include exiting ‘de minimis’ plans, which retained prior non-choosers but could not receive more in

84 exiting plans, with 32,852 beneficiaries being forced to make an active choice or be enrolled in their randomly-selected default plan, 28.3% of whom had previously actively enrolled in their exiting plan.

With this sample, we first show again that overall active choice is still very rare, with only 12.8% of beneficiaries making an active choice to opt out of their randomly-assigned default plan. We then regress an indicator for active choice in the year following exit on an indicator for whether the beneficiary had actively enrolled in her exiting plan. Table 6 presents the results from this regression. In the first column we present the results from a regression with no additional controls. The coefficient represents the difference between the raw means: 7.7% of prior non-active choosers make an active choice, whereas 25.6% of prior active choosers do so. This suggests a role for beneficiary heterogeneity, with prior choice nearly tripling the likelihood of successive active choice.

We next investigate the extent to which this beneficiary heterogeneity is correlated with observable beneficiary characteristics. If the tendency to be an active chooser is not correlated with any observable characteristics, then it will not be possible to target different defaults to active and passive choosers. If that tendency is correlated with observables, on the other hand, default targeting is potentially feasible. In the second column of Table 6, we add in a battery of controls for demographics, including age, gender, and race, as well as controls for health status, i.e. the presence of comorbidities used in the Elixhauser index. Despite these detailed controls, the active choice gap between those who previously made an active choice and those who did not does not decrease. We then add fixed effects for the beneficiary's exiting plan, so that our comparison is done within the plan. While this decreases the effect of prior active choice, the effect still remains. Finally, we allow for the fact that attentiveness may be autocorrelated — i.e. paying attention in one period may encourage the beneficiary to be more lax in the future, or beneficiaries may receive a shock that encourages them to pay attention that slowly decays over time. We do so by interacting the beneficiary's active choice indicator with a measure of the years enrolled in the exiting plan, i.e. the years since her last active choice. The average chooser had been enrolled for 2.9 years, while the average non-chooser had been enrolled for 2.4. We find evidence of decay, with roughly a 4 percentage point decrease in the share of active choice for each year between the present and the prior active choice.

Our interpretation of these results is two-handed. First, these results suggest a real role for permanent heterogeneity in attention, with a significant divide in active choice between prior active beneficiaries and prior passive beneficiaries. Importantly, however, much of this heterogeneity seems to be uncorrelated with observable factors, including the rich demographic and health information we use in our regressions. The fact that rich observable characteristics do not explain the existing heterogeneity suggests that targeting paternalistic defaults to individuals less likely to be “active choice types” is likely to be a fraught exercise, as there do not appear to be any obvious variables to use for targeting.

Second, while active choice in the past does predict active choice today, prior active choice status is not *so* predictive. The modal beneficiary fails to make an active choice regardless of their past status or demographics, and the R^2 of our model is still quite low. This appears to lend credence to the idea that the majority of variation is still from transitory factors.

This latter point should be given some caveats. Interpreting residual variation entirely through such a lens may conflate statistical noise and real transitory variation. Moreover, while this regression models

the year before their exit.

the variation in active choice, it may or may not reflect the variation in latent attentiveness. We therefore conclude with a simple structural exercise that uses parametric assumptions to explicitly estimate the decomposition of variance between permanent and transitory sources of latent attention.

Structural Exercise. Earlier in this section, we noted that $\text{Cov}(A_{it}, A_{it'}) = \text{Var}(c_i)$, but that this was not true for the observed dependent variable a . The analogue, $E[a_{it}|a_{it'}]$, instead depends on the entire distribution of A_{it} . In Appendix D we derive this result explicitly using an application of Bayes' Rule. $E[a_{it}|a_{it'}]$ depends on the marginal distribution of c , $f(c|X)$, where X is a vector of observable individual characteristics, and $E[a_{it'}|c]$, which is a function of the distribution of k . We must make some assumptions about these distributions to proceed. We assume that $c_i \sim \mathcal{N}(\mu X, \sigma^2)$ and $k_{it} \sim \mathcal{N}(0, 1)$. This displays the three potential sources of variation in A_{it} : First, the cross-beneficiary *observable* heterogeneity in X that drives differences in attention, with associated variance $\text{Var}(\mu X)$ (with unknown μ); Second, the cross-beneficiary *unobservable* heterogeneity in c , given by $\text{Var}(c|X)$, which we assume has unknown variance σ^2 ; Finally, the variance in within-beneficiary attention shocks k_{it} across time, which we normalize to 1. We make one further assumption for simplicity: We measure all beneficiary heterogeneity at the time of plan exit and project it back to their initial choice. This assumption will result in our estimation procedure misattributing differences in time-varying observable variation to unobservable transitory shocks, to the extent that our observables vary across time.

Our derivation provides an exact expression for the probability that i makes an active choice in t given their prior active choice in t' and observables X , $P[a_{it} = 1|a_{it'}, X]$, which is computable given our distributional assumptions on c and k . We fit this to our data using maximum likelihood estimation. Our expression for $P[a_{it} = 1|a_{it'}, X]$ has a closed-form expression but no analytic solution; we approximate it using Gauss-Hermite quadrature with 100 nodes. We describe the procedure in detail in Appendix D.

We present the results in Table 7. We estimate one version with no covariates (Column 1) and one with covariates (Column 2). For our covariates we include age (65-74 or 75+), gender, race (white or non-white), and Elixhauser comorbidity count (above or below median). We fully interact these characteristics to get 16 distinct groups and allow each to have a separate mean μ for c . We compute standard errors via bootstrap, using the same bootstrapped sample draws for both models.

In both models, we find that about one-third of the variation in latent attention comes from permanent individual heterogeneity, whereas the other two-thirds comes from transitory shocks. We find that heterogeneity in beneficiary observable features explains only a small amount of the variation in latent attention, about one-tenth of the amount explained by permanent unobservable features. This echoes our prior results: While we observe substantial detail about the beneficiary's health status, which influences the benefit of making a decision, these factors explain hardly any of the variation in active choice. There are many important beneficiary characteristics we do not observe, such as education, (former) occupation, and socioeconomic status. Handel et al. (2020) show that such variables are important in explaining variation in insurance choice mistakes in the Netherlands, and indeed may be the source of this unobserved (in our data) heterogeneity. On the other hand, our results suggest that even a comprehensive dataset richer than what we have here would still be limited in its ability to predict beneficiary default behavior: Two-thirds of the variation occurs *within* beneficiaries across time. We also find that the average beneficiary is fairly inatten-

tive: We estimate that $E[\mu X] = -1$, equal to one standard deviation of the distribution of k . That is, for an average beneficiary to make a choice at all is a 1-sigma event.

Our results in this section suggest that the determinants of active decision-making are not only exogenous to the default given but also largely random. This result bolsters our prior result in supporting generous defaults over shocking ones: Not only must a shocking default overcome the cost of making a choice, it must also overcome random psychological quirks which have an even larger impact. Moreover, this result suggests that, in general, attempts to target behavioral policy towards those most likely to not make an active choice is unlikely to be fruitful, given the inability to observe key sources of heterogeneity and the general randomness that abounds.

8 Conclusion

Defaults are used to steer consumer choices in a wide variety of settings. Recent work in the behavioral economics literature has presented a conflict between those advocating ‘paternalistic’ defaults that steer consumers to products or services that policymakers deem to be better for those consumers (Thaler and Sunstein 2003, Handel and Kolstad 2015) and those advocating ‘shocking’ defaults, designed to induce consumers to make active choices and choose their best options on their own (Choi et al. 2003, Carroll et al. 2009, Beshears et al. forthcoming). Paternalistic defaults are motivated by the idea that consumers are unlikely to ever make an active choice, instead opting to just ‘follow the path of least resistance.’ Shocking defaults, on the other hand, are motivated by the idea that behavioral biases will be overcome if the consequences of passivity are bad enough. In health insurance markets, as well as in other contexts, the optimal default thus depends on the underlying model of consumer choice.

Our results provide strong evidence that shocking defaults are unlikely to be beneficial in the Medicare Part D Low-Income Subsidy program. We show that consumers have a surprisingly strong tendency to follow defaults in this setting, whether the default is to remain enrolled in their incumbent plan or to be forcibly switched to a different, randomly-selected plan. Further, we show almost no relationship between the value of the default and the propensity to make an active choice. These results should cause researchers to exhibit serious caution in advocating for shocking defaults in *any* setting, especially settings similar to the LIS program, such as state Medicaid programs or Affordable Care Act exchanges. Further, our theoretical and empirical results highlight that this relationship between the value of the default and active choice is a critical parameter for default design in any setting. This parameter merits a great deal of future study in a wide variety of settings, and may have implications for other topics, such as the attention response of consumers to changes in product value that come from supplier-driven changes rather than changes in default rules.

Finally, our results suggest that default policies that attempt to target paternalistic defaults to those individuals least likely to make an active choice are unlikely to be an effective default design. This is because none of our rich observables seem to be predictive of heterogeneity in fixed consumer attention “types.” That said, one-third of the variation in active choice does appear to be explained by unobserved fixed individual characteristics. Future research should focus on identifying those characteristics. However, even if those characteristics are identified, a full two-thirds of the variation in active choice remains unexplained, making

it unlikely that targeting defaults will be effective even with a perfect model predicting an individual's attention type. Instead, more work should be dedicated to determining what transitory shocks seem to explain when an individual of a given type decides to make an active choice. Are they random? Or are they a consequence of a specific set of life events? Understanding the answers to these questions presents a fruitful area of future research, with important implications for the design and distributional implications of defaults in health insurance.

References

- Abaluck, Jason and Abi Adams-Prassl**, "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses," 2020.
- **and Jonathan Gruber**, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210.
- Abraham, Sarah and Liyang Sun**, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," 2020.
- Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai**, "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market," *American Economic Review*, forthcoming.
- Bernheim, B. Douglas and Antonio Rangel**, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 51–104.
- **and Dmitry Taubinsky**, "Behavioral Public Economics," in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications*, Vol. 1, North-Holland, 2018, pp. 381 – 516.
- , **Andrey Fradkin, and Igor Popov**, "The Welfare Economics of Default Options in 401(k) Plans," *American Economic Review*, 2015, *105* (9), 2798–2837.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, "The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States," in "Social Security Policy in a Changing Environment," University of Chicago Press, 2009, pp. 167–195.
- , — , — , **and —** , "Active Choice, Implicit Defaults, and the Incentive to Choose," *Organizational Behavior and Human Decision Processes*, forthcoming.
- Blumenstock, Joshua, Michael Callen, and Tarek Ghani**, "Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan," *American Economic Review*, 2018, *108* (10), 2868–2901.
- Bronchetti, Erin T., Judd B. Kessler, Ellen B. Magenheimer, Dmitry Taubinsky, and Eric Zwick**, "Is Attention Produced Rationally? Evidence from Two Experiments," 2020.
- Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad**, "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics," *Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.
- Brown, Zach Y. and Jihye Jeon**, "Endogenous Information Acquisition and Insurance Choice," 2020.

- Caplin, Andrew, Dániel Csaba, John Leahy, and Oded Nov**, “Rational Inattention, Competitive Supply, and Psychometrics,” *Quarterly Journal of Economics*, forthcoming.
- Carroll, Gabriel D, James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick**, “Optimal Defaults and Active Decisions,” *Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight**, “Patient Cost-Sharing and Hospitalization Offsets in the Elderly,” *American Economic Review*, 2010, *100* (1), 193–213.
- Chapman, Gretchen B., Meng Li, and Helen Colby**, “Opting In vs. Opting Out of Influenza Vaccination,” *Journal of the American Medical Association*, 2010, *304* (1), 43–44.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, *99* (4), 1145–1177.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick**, “Defined Contribution Pensions: Plan Rules, Participant Decisions, and the Path of Least Resistance,” *Tax Policy and the Economy*, 2002, *16*, 67–113.
- , —, —, —, and —, “Optimal Defaults,” *American Economic Review*, 2003, *93* (2), 180–185.
- , —, —, —, and —, “For Better or for Worse: Default Effects and 401(k) Savings Behavior,” in “Perspectives on the Economics of Aging,” University of Chicago Press, 2004, pp. 81–126.
- Coughlin, Maura**, “Insurance Choice with Non-Monetary Plan Attributes: Limited Consideration in Medicare Part D,” 2019.
- de Chaisemartin, Clement and Xavier D’Haultfoeuille**, “Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–96.
- DellaVigna, Stefano**, “Structural Behavioral Economics,” in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications*, Vol. 1, North-Holland, 2018, pp. 613 – 723.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf**, “The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D,” *Quarterly Journal of Economics*, 2015, *130* (2), 841–899.
- Ericson, Keith M. Marzilli**, “Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange,” *American Economic Journal: Economic Policy*, February 2014, *6* (1), 38–64.
- Federal Register**, “Patient Protection and Affordable Care Act; HHS Notice of Benefit and Payment Parameters for 2021; Notice Requirement for Non-Federal Governmental Plans,” 2020. Department of Health and Human Services, 45 CFR Parts 146, 149, 155, 156, and 158.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F.P. Luttmer**, “The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment,” *Journal of Political Economy*, 2019, *127* (6), 2836–2874.
- Fowlie, Meredith, Catherine Wolfram, C. Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers**, “Default Effects and Follow-On Behavior: Evidence from an Electricity Pricing Program,” 2017. NBER Working Paper No. 23553.

- Gabaix, Xavier**, “A Sparsity-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 2014, 129 (4), 1661–1710.
- , “Behavioral Inattention,” in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications*, Vol. 2, North-Holland, 2019, pp. 459 – 502.
- Goldin, Jacob and Daniel Reck**, “Optimal Defaults with Normative Ambiguity,” 2020.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” 2019.
- Gruber, Jonathan**, “Delivering Public Health Insurance Through Private Plan Choice in the United States,” *Journal of Economic Perspectives*, 2017, 31 (4), 3–22.
- Handel, Benjamin R.**, “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 2013, 103 (7), 2643–2682.
- and **Jonathan T. Kolstad**, “Getting the Most From Marketplaces: Smart Policies on Health Insurance Choice,” 2015. Brookings Hamilton Project Discussion Paper 2015-08.
- , —, **Thomas Minten**, and **Johannes Spinnewijn**, “The Social Determinants of Choice Quality: Evidence from Health Insurance in the Netherlands,” 2020.
- Heiss, Florian, Adam Leive, Daniel McFadden, Amelie Wupperman, and Joachim Winter**, “Inattention and Switching Costs as Sources of Inertia in Medicare Part D,” 2016. NBER Working Paper No. 22765.
- Ho, Kate, Joseph Hogan, and Fiona Scott Morton**, “The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program,” *RAND Journal of Economics*, 2017, 48 (4), 877–905.
- Hortaçsu, Ali and Chad Syverson**, “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds,” *Quarterly Journal of Economics*, 2004, 119 (2), 403–456.
- Johnson, Eric J. and Daniel Goldstein**, “Do Defaults Save Lives?,” *Science*, 2003, 304 (5649), 1338–1339.
- Keane, Michael P.**, “Modeling Heterogeneity and State Dependence in Consumer Choice Behavior,” *Journal of Business and Economic Statistics*, 1997, 15 (3), 310–327.
- Lee, David S. and David Card**, “Regression Discontinuity Evidence With Specification Error,” *Journal of Econometrics*, 2008, 142 (2), 655–674.
- Madrian, Brigitte C. and Dennis F. Shea**, “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *Quarterly Journal of Economics*, 2001, 116 (4), 1149–1187.
- Matějka, Filip and Alisdair McKay**, “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 2015, 105 (1), 272–298.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, 142 (2), 698–714.
- Morrison, William and Dmitry Taubinsky**, “Rules of Thumb and Attention Elasticities: Evidence from Under- and Overreaction to Taxes,” 2020.
- Polyakova, Maria**, “Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D,” *American Economic Journal: Economic Policy*, 2016, 8 (3), 165–195.

Saez, Emmanuel, “Using Elasticities to Derive Optimal Tax Rates,” *Review of Economic Studies*, 2001, 68 (1), 205–229.


Shampanier, Kristina, Nina Mazar, and Dan Ariely, “Zero as a Special Price: The True Value of Free Products,” *Marketing Science*, 2007, 26 (6), 742–757.

Thaler, Richard H. and Cass R. Sunstein, “Libertarian Paternalism,” *American Economic Review*, 2003, 93 (2), 175–179.


— **and** — , *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Yale University Press, 2008.

Figure 1: Cumulative Active Choice Rate among New Age-65 LIS Beneficiaries

(a) Full U.S. Sample




(b) NY and TX Medicaid-Linked Sample




Notes: Figure plots the cumulative active choice propensity for LIS beneficiaries following their initial Medicare enrollment (at age 65). Panel (a) is based on the sample of all LIS beneficiaries in the U.S. who turned 65 in 2007 - 2010. Panel (b) is based on a subsample of those living in New York and Texas, and who were enrolled in Medicaid at age 64 due to disability before enrolling in Medicare.

Figure 2: Cumulative Active Choice Propensity among New Age-65 LIS Beneficiaries, by Health Status and Value of Choice

(a) By Elixhauser Comorbidity Index, Measured at age 64




(b) By Quartile of Fit for the Beneficiary's Assigned Default Plan




Notes: Figure plots the cumulative active choice propensity for LIS beneficiaries following their initial Medicare enrollment (at age 65), using the New York and Texas Medicaid-Medicare linked sample as in Figure 1b. Panel (a) plots cumulative active choice status by the Elixhauser Comorbidity Index of the beneficiary calculated from their Medicaid inpatient and outpatient claims one year before Medicare enrollment. Panel (b) plots cumulative active choice status by the quartile of fit for the beneficiary's assigned default plan in Medicare Part D, comparing its formulary with prescription drugs consumed at age 64. See Appendix B for more details on the plan fit measure.

Figure 3: Distribution of Difference between Premiums and Market Subsidies




Notes: Figure plots distribution of difference between premiums and market subsidies in year t for plans that were previously benchmark plans in year $t - 1$. The sample includes elderly LIS Medicare beneficiaries who were non active choosers in year $t - 1$ (and therefore were auto-assigned to these benchmark plans). Bins are weighted by prior non-chooser beneficiary enrollment in $t - 1$. Red mass indicates the share of beneficiaries enrolled in 'de minimis' plans that, from 2011 on, were allowed to retain their incumbent beneficiaries. See Section 3.1 for more discussion of how these plans are treated in analyses.

Figure 4: Regression Discontinuity Balance Tests for Observable Beneficiary Characteristics




Notes: Figure plots the means of observed beneficiary characteristics for elderly LIS Medicare beneficiaries in year $t - 1$ (prior to the change in the default from remaining in the beneficiary’s incumbent plan to being auto-enrolled in a randomly-selected default plan), for whole-dollar bins of the difference between their incumbent plans’ premium in year t and the regional subsidy in that year. Panels (a) and (b) are based on the full national sample, excluding “de minimis” plans. Panels (c) - (e) are based on a 20 percent random subsample, for whom we observe inpatient, outpatient, and prescription drug claims, also excluding “de minimis” plans. Panel (e) plots the average number of Elixhauser Comorbidities for beneficiaries in each bin.

Figure 5: Probability of Switching Plans, by Incumbent Plan Premium Relative to Benchmark




Notes: Figure plots the share of elderly LIS beneficiaries who enroll in different plans in January of year t versus December of year $t - 1$, for whole-dollar bins of the difference between their incumbent plans' premium in year t and the regional subsidy in that year. Plots are based on the full national sample, excluding "de minimis" plans.

Figure 6a: Theoretical Relationship Between ‘Stakes’ and Active Choice




Notes: This figure plots the relationship between potential welfare loss from failing to make an active choice (driven by variation in the value of the default), and the ex ante probability of making an active choice under various models of active decision-making. The blue line represents the ‘rational inattention’ case where beneficiaries put equal weight on the costs of making an active choice and the expected benefits of doing so, $\beta = 1$ and $k = 0$. The green line represents the ‘boundedly-rational inattention’ case where beneficiaries ‘over-weight’ costs relative to benefits, with $0 < \beta < 1$ and $k = 0$. The brown line represents the ‘random inattention’ case where beneficiaries randomly make active choices, with $\beta = c = 0$ and $k \sim U[c - (1 - \theta), c + \theta]$.

Figure 6b: Theoretical Relationship Between ‘Stakes’ and Welfare Loss




Notes: This figure plots the relationship between potential welfare loss from failing to make an active choice (driven by variation in the value of the default), and the realized loss (taking the probability of active choice into account) under various models of active decision-making. The blue line represents the ‘rational inattention’ case where beneficiaries put equal weight on the costs of making an active choice and the expected benefits of doing so, $\beta = 1$ and $k = 0$. The green line represents the ‘boundedly-rational inattention’ case where beneficiaries ‘over-weight’ costs relative to benefits, with $0 < \beta < 1$ and $k = 0$. The brown line represents the ‘random inattention’ case where beneficiaries randomly make active choices, with $\beta = c = 0$ and $k \sim U[c - (1 - \theta), c + \theta]$.

Figure 7: Welfare Effects of Default Changes




Notes: This figure plots the relationship between potential welfare loss from failing to make an active choice (driven by the value of the default) versus the realized loss (taking into account the probability of making an active choice) for a boundedly rational consumer with $0 < \beta < 1$ and $k = 0$. The magenta line plots the same relation when v^d is changed by $\Delta v^d < 0$, with the x-axis representing potential welfare loss under the initial default, $v^* - v^d$, such that ticks on the axis represent the same individuals across default regimes. The dashed section of the magenta line represents the welfare loss avoided by marginal beneficiaries induced to make an active choice.

Figure 8: Event Study – Effect of Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Chosen Default Plan on Drug Consumption




Notes: This figure plots event-study regression estimates of the effect of a change in the default from remaining in the incumbent plan to auto-assignment to a randomly-chosen default plan on drug consumption. We use a stacked difference-in-differences design (described in detail in Section 5) to compare outcomes pre- vs. post-default change for beneficiaries whose year $t - 1$ plans bid up to \$6 above the year t benchmark (treatment group) vs. beneficiaries whose year $t - 1$ plans bid up to \$6 below the year t benchmark (control group). For this, the RD analytic sample is employed, which is described in further detail in Table 1. Each market-year represents its own difference-in-differences ‘experiment.’ The regression estimates the average effect across all experiments. Regressions include beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and quarter indicators relative to experiment time. Standard errors are clustered at the level of the beneficiary. The outcome variable in panel (a) is quarterly prescription drug spending, and the outcome in panel (b) is price-normalized quarterly prescription drug spending, where prices are normalized to the average price (at the NDC level) across all plans in the sample for that year.

Figure 9: Event Study – Effect of Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Chosen Default Plan on Drug Consumption, by Fit of New Default Plan




Notes: This figure plots event-study regression estimates of the effect of a change in the default from remaining in the incumbent plan to auto-assignment to a randomly-chosen default plan on drug consumption. Blue dots (red triangles) represent effects for beneficiaries assigned default plans in the bottom quintile (top four quintiles) of beneficiary-specific ranking of plan fit. We use a stacked difference-in-differences design (described in detail in Section 5) to compare outcomes pre- vs. post-default change for beneficiaries whose year $t - 1$ plans bid up to \$6 above the year t benchmark (treatment group) vs. beneficiaries whose year $t - 1$ plans bid up to \$6 below the year t benchmark (control group). For this, the RD analytic sample is employed, which is described in further detail in Table 1. Each market-year represents its own difference-in-differences ‘experiment.’ The regression estimates the average effect across all experiments. Regressions include beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and quarter indicators relative to experiment time. Standard errors are clustered at the level of the beneficiary. The outcome variable in panel (a) is quarterly prescription drug spending, and the outcome in panel (b) is price-normalized quarterly prescription drug spending, where prices are normalized to the average price (at the NDC level) across all plans in the sample for that year.

Figure 10: Event Study – Effect of Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Chosen Default Plan on Switching and Active Choice, by Fit of New Default Plan



Notes: This figure plots event-study regression estimates of the effect of a change in the default from remaining in the incumbent plan to auto-assignment to a randomly-chosen default plan on switching and active choice outcomes. Blue dots (red triangles) represent effects for beneficiaries assigned default plans in the bottom quintile (top four quintiles) of beneficiary-specific ranking of plan fit. We use a stacked difference-in-differences design (described in detail in Section 5) to compare outcomes pre- vs. post-default change for beneficiaries whose year $t - 1$ plans bid up to \$6 above the year t benchmark (treatment group) vs. beneficiaries whose year $t - 1$ plans bid up to \$6 below the year t benchmark (control group). For this, the RD analytic sample is employed, which is described in further detail in Table 1. Each market-year represents its own difference-in-differences ‘experiment.’ The regression estimates the average effect across all experiments. Regressions include Part D region-experiment year fixed effects and quarter indicators relative to experiment time. Standard errors are clustered at the level of the beneficiary. The outcome variable in panel (a) is a binary indicator for whether a beneficiary is in a different plan in December of year t than they were in December of year $t - 1$. The outcome in panel (b) is a binary indicator for whether a beneficiary had made an active choice to enroll in the plan they were in as of December of year t .

Figure 11: Comparison of the Reduction in Consumption Due to the Change in Default Versus Effects of the Change in Default in Switching and Active Choice



Notes: The top panel plots the propensity of beneficiaries to switch from the plan they were enrolled in during December of year $t - 1$, as a function of the potential consumption loss they face by remaining in their default plan assignment over year t , in terms of log drug spending. The bottom panel plots the propensity of beneficiaries to make an active choice in December of year t , as a function of the potential consumption loss they face by remaining in their default plan assignment. In both panels, estimates of consumption loss come from a regression of log drug spending on whether a beneficiary's plan loses benchmark status in year t , and includes beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and experiment pre/post year indicators. Estimates of propensity to switch and propensity to make an active choice come from regressions also including whether a beneficiary's plan lost benchmark status in year t , as well as region-experiment year fixed effects. Each regression is run for five different treatment groups, based on quintile of assigned plan fit, while full control group in all regressions. Matching coefficients from each of these five regression specifications are then plotted in the figures.

Table 1: Summary statistics for samples used in our analysis.

	Broad LIS Sample	National MCR Initial Enrollment Sample	NY TX MCD-MCR Linked Sample	RD Analysis Sample
Avg. Age	77.2	66	66	76.6
% Female	69.5%	62.2%	62.2%	69.1%
Avg. Elixhauser Index	4.5	3.4	4.4	4.2
Total Drug Spending				
Mean	\$4,024	\$3,007	\$4,231	\$3,707
25th Percentile	\$815	\$298	\$871	\$710
Median	\$2,601	\$1,561	\$2,661	\$2,384
75th Percentile	\$5,389	\$3,945	\$5,364	\$5,021
Avg. Chronic Drug Spending	\$3,012	\$2,264	\$3,230	\$2,821
Avg. Non-Drug Med. Spending	\$12,888	\$8,413	\$11,329	\$10,945
Observations	4,628,704	216,772	14,218	1,989,603

Notes: The table includes summary statistics for four samples. For all of our analyses in this paper, we employ a set of common restrictions in the samples. We restrict only to dual-eligibles (those simultaneously enrolled in Medicaid and Medicare, and thus receiving full LIS benefits) over 65 enrolled in Medicare Parts A, B, and D, and not enrolled in Medicare Advantage. The “Broad LIS Sample” is a random 20% subset of that sample, for whom we can observe detailed medical and drug utilization in Medicare. The “National MCR Initial Enrollment Sample” restricts to the first year of beneficiaries’ Medicare enrollment (at age 65). The “NY TX MCD-Linked Sample” further restricts to beneficiaries residing in New York and Texas who were on Medicaid due to disability before enrolling in Medicare at age 65. The “RD Analysis Sample” includes (20 % subset of) LIS beneficiaries who were enrolled in a benchmark Part D plan (in year t-1) and continued to enroll in Medicare in the following year (t). The Elixhauser Index is a count of co-morbidities, calculated from all Medicare inpatient, outpatient, and carrier files in the corresponding beneficiary-year for the “Broad LIS Sample”, in first calendar year post initial Medicare enrollment for the “National MCR Initial Enrollment Sample” and “NY TX MCD-Linked Sample”, and in year t-1 for the “RD Analysis Sample”. Meanwhile, the spending measures are based on Medicare data in the corresponding beneficiary-year for the “Broad LIS Sample”, in first calendar year post initial Medicare enrollment for the “National MCR Initial Enrollment Sample” and “NY TX MCD-Linked Sample”, and in year t for the “RD Analysis Sample”. Chronic drugs are defined as NDC codes where the median number of fills per person in a year is greater than two, among the subset of people with at least one fill.

Table 2: Regression Discontinuity Estimates of the Effects of a Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Selected Default Plan

	Switch in January of t			
	Overall	By Active Choice	By Active Choice to Benchmark Plan	By Default Reassignment
$1(Bid_{j(i)t} - Benchmark_{r(i)t})$	0.960 (0.001)	0.024 (0.001)	0.020 (0.001)	0.936 (0.001)
Average Value for Control Beneficiaries	0.006	0.006	0.004	0
Observations	460,729			

Notes: This table reports regression discontinuity estimates of the effect of a change in a beneficiary’s default from remaining in their incumbent plan to auto-assignment to a randomly-selected default plan on their probability of switching plans and active choice. The table reports coefficients on a dummy variable for the beneficiary’s year $t - 1$ plan j having a year t bid that exceeds the year t benchmark in region r . For this, the RD analytic sample is employed, which is described in further detail in Table 1. RD regressions include experiment fixed effects, where an ‘experiment’ is a Part D region-year pair. All regressions restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Column (1) provides estimates of effects on the probability of enrolling in a different plan in year t than in year $t - 1$. Column (2) provides estimates of effects on the probability of making an active choice in year t . Column (3) provides estimates of the effects on the probability of actively choosing a benchmark plan (as opposed to a non-benchmark plan). Column (4) provides estimates of effects on the probability of switching due to default re-assignment (rather than active choice).

Table 3: Difference-in-Difference Regression Estimates of the Effects of a Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Selected Default Plan on Drug Consumption

Price Normalization	Log Drug Spending		
	-	Drug	Class
Panel A			
$1(Bid_{j(i)t} - Benchmark_{r(i)t})$ × Post	-0.064 (0.004)	-0.050 (0.004)	-0.021 (0.004)
Panel B			
$1(Bid_{j(i)t} - Benchmark_{r(i)t})$ × Post	-0.043 (0.004)	-0.030 (0.004)	-0.012 (0.004)
$1(Bid_{j(i)t} - Benchmark_{r(i)t})$ × Bottom Quintile Fit × Post	-0.083 (0.007)	-0.082 (0.007)	-0.036 (0.006)
Average Value for Control Beneficiaries in $t - 1$ (levels)	\$3,329	\$3,412	\$3,503
Observations	5,574,684		

Notes: This table reports difference-in-difference regression estimates of the change in a beneficiary’s default from remaining in their incumbent plan to auto-assignment to a randomly-selected default plan on yearly drug spending outcomes. The treatment group consists of beneficiaries whose year $t - 1$ plans bid above the year t benchmark, and the control group consists of beneficiaries whose year $t - 1$ plans j bid below the year t benchmark in region r . For this, the RD analytic sample is employed, which is described in further detail in Table 1. For both groups, we restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Regressions include beneficiary-experiment level fixed effects and experiment-pre/post year indicators, where an ‘experiment’ is a Part D region and year pair. Standard errors are clustered at the level of the beneficiary. In Panel A, we estimate the overall effect of the change in default. In Panel B, we estimate heterogeneous effects by the fit of the plan (randomly) selected to be the beneficiary’s year t default by interacting an indicator for being assigned a default plan in the bottom quintile of fit among the benchmark plans in the beneficiary’s choice set with dummies for “Post” and $1(Bid_t > Benchmark_t)$. As such, the coefficients in the first row of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the top four quintiles of fit, and the sum of the coefficients in the first and second rows of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the bottom quintile of fit. In each column, we apply a different normalization of drug prices. The first column has no normalization. The second normalizes prices at the drug (NDC) level. The third normalizes prices at the drug class level. In both cases, we normalize by finding the average price across all plans and assigning that price to all relevant drugs.

Table 4: Difference-in-Difference Regression Estimates of the Effects of a Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Selected Default Plan on Additional Drug Consumption Outcomes

	Log Drug Spending		
	High-Value	Chronic	Non-Chronic
Panel A			
Benchmark Status Loss × Post	-0.062 (0.005)	-0.078 (0.005)	-0.028 (0.006)
Panel B			
Benchmark Status Loss × Post	-0.049 (0.005)	-0.057 (0.005)	-0.022 (0.006)
Benchmark Status Loss × Bottom Quintile Fit × Post	-0.052 (0.007)	-0.085 (0.008)	-0.023 (0.009)
Average Value for Control Beneficiaries in $t - 1$ (levels)	\$1,237	\$2,496	\$834
Observations	5,574,684		

Notes: This table reports difference-in-difference regression estimates of the change in a beneficiary’s default from remaining in their incumbent plan to auto-assignment to a randomly-selected default plan on yearly drug spending outcomes. The treatment group consists of beneficiaries whose year $t - 1$ plans bid above the year t benchmark, and the control group consists of beneficiaries whose year $t - 1$ plans bid below the year t benchmark. For this, the RD analytic sample is employed, which is described in further detail in Table 1. For both groups, we restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Regressions include beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and experiment-pre/post year indicators. Standard errors are clustered at the level of the beneficiary. In Panel A, we estimate the overall effect of the change in default. In Panel B, we estimate heterogeneous effects by the fit of the plan (randomly) selected to be the beneficiary’s year t default by interacting an indicator for being assigned a default plan in the bottom quintile of fit among the benchmark plans in the beneficiary’s choice set with dummies for “Post” and $1(Bid_t > Benchmark_t)$. As such, the coefficients in the first row of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the top four quintiles of fit, and the sum of the coefficients in the first and second rows of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the bottom quintile of fit. In column (1), the outcome is spending on “high-value” drugs, which are defined based on whether they belong to therapeutic classes targeting key chronic conditions including diabetes, depression, hypertension, and asthma. In columns (2) and (3), the outcomes are spending on chronic and non-chronic drugs. Chronic drugs are defined as drugs for which the median number of annual, per person fills is greater than two, among the subset of people with at least one annual fill. Non-chronic drugs are any drugs not classified as chronic drugs.

Table 5: Difference-in-Difference Regression Estimates of the Effects of a Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Selected Default Plan on Switching and Active Choice

	Switch From $t - 1$ Plan by December of t	Active Choice by December of t
Panel A		
Benchmark Status Loss	0.955 (0.001)	0.050 (0.001)
Panel B		
Benchmark Status Loss	0.955 (0.001)	0.047 (0.001)
Benchmark Status Loss × Bottom Quintile Fit	0.002 (0.002)	0.015 (0.002)
Average Value for Control Beneficiaries	0.011	0.010
Observations	460,729	

Notes: This table reports difference-in-difference regression estimates of the change in a beneficiary’s default from remaining in their incumbent plan to auto-assignment to a randomly-selected default plan on switching and active choice outcomes. The treatment group consists of beneficiaries whose year $t - 1$ plans bid above the year t benchmark, and the control group consists of beneficiaries whose year $t - 1$ plans bid below the year t benchmark. For both groups, we restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Regressions include Part D region-experiment-year fixed effects. In Panel A, we estimate the overall effect of the change in default. In Panel B, we estimate heterogeneous effects by the fit of the plan (randomly) selected to be the beneficiary’s year t default by interacting an indicator for being assigned a default plan in the bottom quintile of fit among the benchmark plans in the beneficiary’s choice set with dummies for "Post" and $1(Bid_t > Benchmark_t)$. As such, the coefficients in the first row of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the top four quintiles of fit, and the sum of the coefficients in the first and second rows of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the bottom quintile of fit. In column (1), the outcome is a binary indicator for whether a beneficiary is in a different plan in December of year t than they were in December of year $t - 1$. In column (2), the outcome is a binary indicator for whether a beneficiary had made an active choice to enroll in the plan they were in as of December of year t .

Table 6: Estimates of Effects of Prior Active Choice and Other Observed and Unobserved Individual Characteristics on Current Active Choice

	Made Active Choice Following Plan Exit			
	(1)	(2)	(3)	(4)
Actively Enrolled in Exiting Plan	0.179 (0.005)	0.175 (0.005)	0.119 (0.006)	0.237 (0.011)
Years Enrolled in Exiting Plan				0.011 (0.003)
Actively Enrolled \times Years Enrolled				-0.038 (0.002)
Demographic and Health Controls		X	X	X
Year Fixed Effects		X	X	X
Exiting Plan Fixed Effects			X	X
R^2	0.058	0.092	0.181	0.190
Share Active Choice			0.128	
Share Active Choice For Previous Non-Active			0.077	
Share Active Choice For Previous Active			0.256	
Share Previous Active			0.283	
Observations			32,852	
Number of Exiting Plans			84	

Notes: This table reports estimates from linear regressions of a beneficiary's active choice propensity following the exit of their incumbent plan on whether they had previously actively elected to enroll in the that plan. The analysis sample includes LIS beneficiaries aged over 65 who in year $t - 1$ were enrolled in a benchmark plan (including "de minimis" plans) which exited the market in year t . In the first column we present the results from a regression with no additional controls. In column (2), we add in a battery of controls for demographics, including age, gender, and race, as well as controls for health status, i.e. the presence of comorbidities used in the Elixhauser Index. In column (3), we add fixed effects for the beneficiary's exiting plan, so that our comparison is done within the plan. Lastly, to allow for the fact that attentiveness may be autocorrelated, we interact the beneficiary's active choice indicator with a measure of the years enrolled in the exiting plan, i.e., the years since their last active choice, in column (4).

Table 7: Structural Model of Active Choice Parameter Estimates


	(1)	(2)
$E[\mu X]$	-1.01 [-1.03,-0.99]	-1.01 [-1.03, -0.99]
σ	0.71 [0.68,0.74]	0.68 [0.65,0.71]
Share of Variance		
Var $[\mu X_i]$ (Observable Permanent Variation)	-	3.2% [2.9%, 3.8%]
Var $[c_i X_i] = \sigma^2$ (Unobservable Permanent Variation)	33.4% [31.4%, 35.4%]	30.6% [28.7%, 32.5%]
Var $[k_{it}]$ (Unobservable Transitory Variation)	66.6% [64.6%, 68.6%]	66.2% [64.1%, 68.0%]
Observations	32,852	

Notes: The upper panel of this table reports estimates produced from our structural parametric model of attention following plan exit. The lower panel details the share of variance in latent attention described by the associated factor, derived from the parameter estimates of the models. Below the estimate we give the 95% confidence interval generated from 1000 bootstrap runs. Observable characteristics in X_i include age, gender, and race, as well as controls for health status, i.e. the presence of comorbidities used in the Elixhauser Index.


For Online Publication

Appendix for:
**The Behavioral Foundations of Default Effects:
Theory and Evidence from Medicare Part D**


A Appendix: Additional Figures and Tables




Appendix Figure A1: Cumulative active choice propensity for our Medicaid-linked sample by number of unique drugs taken at 64.




Appendix Figure A2: This figure plots premiums in $t - 1$ against premiums in t for all plans in our data that existed in consecutive year pairs.




Appendix Figure A3: This figure replicates Figure 3, restricting only to 2007-2010, before the de minimis rules were implemented.




Appendix Figure A4: This figure replicates Figure 3, restricting only to 2011 onwards, after the de minimis rules were implemented.




Appendix Figure A5: We break down the outcome in Figure 5 by whether the plan switch was an active choice, an active choice specifically to a non-benchmark plan, or was a reassignment through the default mechanism.



Appendix Figure A6: This figure replicates Figure 5, but additionally plots switching rates among those who had actively chosen their incumbent plan.




Appendix Figure A7: Probability of switching plans between December of $t - 1$ and December of t , by incumbent plan premium in t .




Appendix Figure A8: This figure plots the relationship between potential welfare loss from failing to make an active choice, and ex ante probability of active choice. Each line represents an agent following a model of rational decision costs with $\beta = 1$ and $k = 0$, with decision costs c distributed randomly in $U[\bar{c} - \epsilon, \bar{c} + \epsilon]$. The blue line represents $\epsilon = 0$, while redder lines represent higher values of ϵ , in $\{0.5\bar{c}, \bar{c}, 2\bar{c}, 3\bar{c}, 4\bar{c}, 8\bar{c}\}$.


Appendix Figure A9: Distribution of Number of Benchmark Plans in Market-Year



Notes: This set of figures plots the distribution in the number of benchmark plans across the combination of Part D market region-years. The top figure presents this distribution weighing all Part D market region-years equally, while the bottom weighs Part D market region-years by the number of beneficiaries in our sample enrolled under each.

Appendix Figure A10: Rational Inattention, Top Half in Spread of Fit

Notes: This figure replicates Figure 11, restricting to beneficiaries for whom the gap between the fit of their best plan and their worst plan is above the median gap in the sample. The top panel plots the propensity of beneficiaries to switch from the plan they were enrolled in during December of year $t - 1$, as a function of the potential consumption loss they face by remaining in their default plan assignment over year t , in terms of log drug spending. The bottom panel plots the propensity of beneficiaries to make an active choice in December of year t , as a function of the potential consumption loss they face by remaining in their default plan assignment. In both panels, estimates of consumption loss come from a regression of log drug spending on whether a beneficiary's plan loses benchmark status in year t , and includes beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and experiment pre/post year indicators. Estimates of propensity to switch and propensity to make an active choice come from regressions also including whether a beneficiary's plan lost benchmark status in year t , as well as region-experiment year fixed effects. Each regression is run for five different treatment groups, based on quintile of assigned plan fit, while full control group in all regressions. Matching coefficients from each of these five regression specifications are then plotted in the figures.

Appendix Figure A11: Rational Inattention, Top Quarter in Spread of Fit

Notes: This figure replicates Figure 11, restricting to beneficiaries for whom the gap between the fit of their best plan and their worst plan is above the 75th percentile gap in the sample. The top panel plots the propensity of beneficiaries to switch from the plan they were enrolled in during December of year $t - 1$, as a function of the potential consumption loss they face by remaining in their default plan assignment over year t , in terms of log drug spending. The bottom panel plots the propensity of beneficiaries to make an active choice in December of year t , as a function of the potential consumption loss they face by remaining in their default plan assignment. In both panels, estimates of consumption loss come from a regression of log drug spending on whether a beneficiary's plan loses benchmark status in year t , and includes beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and experiment pre/post year indicators. Estimates of propensity to switch and propensity to make an active choice come from regressions also including whether a beneficiary's plan lost benchmark status in year t , as well as region-experiment year fixed effects. Each regression is run for five different treatment groups, based on quintile of assigned plan fit, while full control group in all regressions. Matching coefficients from each of these five regression specifications are then plotted in the figures.

Appendix Table A1: Regression Discontinuity Balance Tests for Observable Beneficiary Characteristics

	Age	% Female	Prior Drug Spending	Prior Non-Drug Medical Spending	Elixhauser Comorbidities
$1(Bid_t > Benchmark_t)$	0.021 (0.034)	0 (0.002)	-23.484 (19.283)	-23.102 (85.557)	-0.001 (0.013)
Average Value for Control Beneficiaries	76.3	0.682	3,328	9,528	3.88
Observations	464,341				

Notes: This table reports regression discontinuity balance tests for observable beneficiary characteristics, illustrated by coefficients on a dummy variable for the beneficiary's year $t - 1$ plan having a year t bid that exceeds the year t benchmark. For this, the RD analytic sample is employed, which is described in further detail in Table 1. RD regressions include Part D region-experiment year fixed effects. All regressions restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Regression coefficients summarize the results in Figure 4.

Appendix Table A2: Difference-in-Difference Regression Estimates of the Effects of a Change in Default from Remaining in Incumbent Plan to Auto-assignment to Randomly-Selected Default Plan on Drug and Non-Drug Outcomes

	Log # Fills	Log Days Supply	Log Non-Drug Spending
Panel A			
Benchmark Status Loss × Post	-0.005 (0.002)	-0.015 (0.003)	0.006 (0.006)
Panel B			
Benchmark Status Loss × Post	-0.002 (0.002)	-0.007 (0.004)	0.010 (0.007)
Benchmark Status Loss × Bottom Quintile Fit × Post	-0.011 (0.003)	-0.030 (0.006)	-0.017 (0.010)
Average Value for Control Beneficiaries in $t - 1$ (levels)	53.3	1649	\$9535
Observations		5,574,684	

Notes: This table reports difference-in-difference regression estimates of the change in a beneficiary's default from remaining in their incumbent plan to auto-assignment to a randomly-selected default plan on yearly drug utilization and non-drug spend outcomes. The treatment group consists of beneficiaries whose year $t - 1$ plans bid above the year t benchmark, and the control group consists of beneficiaries whose year $t - 1$ plans bid below the year t benchmark. For this, the RD analytic sample is employed, which is described in further detail elsewhere. For both groups, we restrict to beneficiaries whose year $t - 1$ plans bid within \$6 of the year t benchmark. Regressions include beneficiary-experiment level fixed effects, Part D region-experiment year fixed effects, and experiment-pre/post year indicators. In Panel A, we estimate the overall effect of the change in default. In Panel B, we estimate heterogeneous effects by the fit of the plan (randomly) selected to be the beneficiary's year t default by interacting an indicator for being assigned a default plan in the bottom quintile of fit among the benchmark plans in the beneficiary's choice set with dummies for "Post" and $1(Bid_t > Benchmark_t)$. As such, the coefficients in the first row of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the top four quintiles of fit, and the sum of the coefficients in the first and second rows of Panel B can be interpreted as the effects of the change in default for beneficiaries whose randomly-selected default plan is in the bottom quintile of fit. In column (1), the outcome is number of prescription fills, while the outcome in column (2) is total days supply associated with those fills. In column (3) meanwhile, the outcome is non-drug medical spending, defined as spending under Medicare Parts A and B.

Appendix Table A3: Event Study — Effect of Change in Default from Remaining in Incumbent Plan to Auto- assignment to Randomly-Chosen Default Plan

	Switch from $t - 1$ Plan by December of t	Active Choice by December of t	Log (Drug Spending)	Log (Price-Normalized Drug Spending)
T - 8			0.013 (0.007)	0.021 (0.007)
T - 7			0.001 (0.007)	0.015 (0.007)
T - 6			0.009 (0.006)	0.022 (0.006)
T - 5			0.009 (0.006)	0.022 (0.006)
T - 4			0.011 (0.006)	0.014 (0.006)
T - 3			-0.003 (0.005)	-0.005 (0.005)
T - 2			0.004 (0.005)	-0.001 (0.005)
T	0.960 (0.001)	0.046 (0.001)	-0.041 (0.005)	-0.037 (0.005)
T + 1	0.957 (0.001)	0.049 (0.001)	-0.077 (0.006)	-0.063 (0.005)
T + 2	0.956 (0.001)	0.050 (0.001)	-0.075 (0.006)	-0.054 (0.006)
T + 3	0.955 (0.001)	0.050 (0.001)	-0.050 (0.006)	-0.038 (0.006)
Mean of Dep Var	0.170	0.016	5.390	5.451
Observations	2,307,468	2,307,441	5,574,672	5,574,672

Notes: Regression results underlying Figure 8

Appendix Table A4: Event Study — Effect of Change in Default from Remaining in Incumbent Plan to Auto- assignment to Randomly-Chosen Default Plan

	Switch from $t - 1$ Plan by December of t	Active Choice by December of t	Log (Drug Spending)	Log (Price-Normalized Drug Spending)
T - 8 (low fit)			0.018 (0.011)	0.026 (0.011)
T - 8 (high fit)			0.011 (0.007)	0.020 (0.007)
T - 7 (low fit)			-0.003 (0.010)	0.011 (0.010)
T - 7 (high fit)			0.002 (0.007)	0.017 (0.007)
T - 6 (low fit)			0.006 (0.010)	0.019 (0.010)
T - 6 (high fit)			0.009 (0.007)	0.023 (0.007)
T - 5 (low fit)			0.011 (0.009)	0.022 (0.009)
T - 5 (high fit)			0.009 (0.006)	0.022 (0.006)
T - 4 (low fit)			0.015 (0.009)	0.017 (0.009)
T - 4 (high fit)			0.009 (0.006)	0.013 (0.006)
T - 3 (low fit)			-0.002 (0.008)	-0.004 (0.008)
T - 3 (high fit)			-0.004 (0.006)	-0.005 (0.006)
T - 2 (low fit)			0.012 (0.008)	0.007 (0.008)
T - 2 (high fit)			0.001 (0.005)	-0.003 (0.005)
T (low fit)	0.962 (0.001)	0.055 (0.002)	-0.071 (0.008)	-0.067 (0.008)
T (high fit)	0.959 (0.001)	0.043 (0.001)	-0.031 (0.006)	-0.027 (0.006)
T + 1 (low fit)	0.959 (0.001)	0.059 (0.002)	-0.159 (0.009)	-0.144 (0.009)
T + 1 (high fit)	0.957 (0.001)	0.046 (0.001)	-0.050 (0.006)	-0.036 (0.006)
T + 2 (low fit)	0.957 (0.001)	0.061 (0.002)	-0.142 (0.009)	-0.118 (0.009)
T + 2 (high fit)	0.956 (0.001)	0.047 (0.001)	-0.053 (0.006)	-0.033 (0.006)
T + 3 (low fit)	0.957 (0.001)	0.062 (0.002)	-0.109 (0.009)	-0.100 (0.009)
T + 3 (high fit)	0.955 (0.001)	0.047 (0.001)	-0.030 (0.006)	-0.017 (0.006)
Mean of Dep Var	0.170	0.016	5.390	5.451
Observations	2,307,468	2,307,441	5,574,672	5,574,672

Notes: Regression results underlying Figures 9a, 9b, 10a, 10b.

Appendix Table A5: Comparison of the Reduction in Consumption Due to the Change in Default Versus Effects of the Change in Default in Switching and Active Choice

	Fit Quintile				
	1	2	3	4	5
	Original				
Active Choice Response	0.073 (0.002)	0.063 (0.002)	0.054 (0.002)	0.051 (0.002)	0.05 (0.002)
Switch Response	0.954 (0.001)	0.951 (0.001)	0.957 (0.001)	0.955 (0.001)	0.953 (0.002)
Drug Response	-0.132 (0.005)	-0.082 (0.006)	-0.063 (0.006)	-0.055 (0.006)	-0.025 (0.006)
	TOP 50% Spread				
Active Choice Response	0.088 (0.003)	0.07 (0.003)	0.055 (0.003)	0.051 (0.002)	0.045 (0.003)
Switch Response	0.956 (0.002)	0.959 (0.002)	0.953 (0.002)	0.946 (0.002)	0.942 (0.002)
Drug Response	-0.241 (0.009)	-0.162 (0.010)	-0.081 (0.010)	-0.033 (0.010)	-0.008 (0.011)
	TOP 25% Spread				
Active Choice Response	0.083 (0.004)	0.06 (0.004)	0.056 (0.004)	0.047 (0.003)	0.04 (0.003)
Switch Response	0.958 (0.002)	0.961 (0.003)	0.962 (0.002)	0.958 (0.003)	0.932 (0.003)
Drug Response	-0.271 (0.014)	-0.147 (0.015)	-0.085 (0.015)	-0.096 (0.016)	0.048 (0.015)

Notes: Regression results underlying Figures 11a, 11b.

B Appendix: Measurement of Plan ‘Fit’


In this appendix, we describe our measure of how well a plan’s formulary “fits” the basket of drugs demanded by a beneficiary. For each beneficiary, we identify the basket of drugs taken by the beneficiary in year $t - 1$. For new beneficiaries, $t - 1$ is the year prior to the year the beneficiary entered Medicare, and the basket of drugs is constructed using data from the Medicaid programs in New York and Texas in the year prior to entering Medicare at age 65. For beneficiaries experiencing a change in their default from staying in their incumbent plan to auto-assignment to a randomly-selected default plan, $t - 1$ is the year prior to the change in the default and the basket of drugs is constructed using data from that year (when beneficiaries were in their prior plan). The measure ranges from 0 (no drugs in the beneficiary’s basket covered by the plan) to 1 (all drugs in the beneficiary’s basket covered by the plan).

We plot the distribution of these beneficiary-plan fit measures for benchmark plans in Figure A12. Plan-beneficiary fit varies significantly for beneficiaries in the LIS program – while some plans cover 100% of a beneficiary’s drugs, others cover none at all, with nearly full support across the distribution between 0 and 1. The average benchmark plan covers 60% of a beneficiary’s prescriptions with a standard deviation of 23 percentage points.


Ultimately, we use the measure of plan-beneficiary fit to (1) measure the value of active choice and (2) assess the consequences of being assigned to a low-fitting plan versus a high-fitting plan. Both of these uses depend more on a beneficiary facing *differences* in fit across plans rather than lower or higher average fit across plans. Because of this, we present information on the *within-beneficiary* variation in fit across plans in the beneficiary’s choice set. Figure A13 plots, for each beneficiary, the difference between the fit of the best fitting plan and the average fit across plans in the beneficiary’s choice set.³³ This distribution is more compressed than the distribution of fit, owing to the correlation between fit values across plans for a given beneficiary. Despite this, we still see a great deal of heterogeneity, with a significant portion of beneficiaries facing relative coverage gaps of over 10 percentage points and some beneficiaries facing relative coverage gaps nearing half of their drugs.

In Figure A14, we create a plot similar to Figure 2b, but using quartiles of the raw difference between the best-fitting plan a beneficiary has access to and the average, another representation of the “value of choice.” In Section 6 we explain why using absolute differences may not be ideal. However, it produces results similar to that of Figure 2b.


³³One important note is that we take the maximum fit value from only the set of benchmark plans. Some beneficiaries might be able to improve their coverage by choosing a non-free insurance option. In practice, few LIS beneficiaries end up in such plans. Our approach is agnostic about the beneficiary’s preferred trade-off between formulary coverage and premium payments and therefore serves as a lower bound on the true value of choice.



Appendix Figure A12: Distribution of beneficiary-plan pairwise fit measures for our Medicaid-linked sample of 65-year-olds, where fit is defined as the share of drugs taken at 64 covered by the plan. Fit is only computed for benchmark plans available when they turned 65.



Appendix Figure A13: Distribution of beneficiary-specific measures of the difference between the fit of the best-fitting plan, and the average across available benchmark plans when they turned 65, for our Medicaid-linked sample of 65-year olds.



Appendix Figure A14: Cumulative active choice propensity for our Medicaid-linked sample by quartile of our 'value of choice' measure.


C Appendix: Long-Run Effects

In Sections 5 and 6, we use a regression discontinuity design leveraging the fact that when a $t - 1$ benchmark plan submits a year t bid exceeding the year t subsidy, the beneficiaries in that plan face an exogenous change in their default from remaining in their year t plan to auto-assignment to a randomly-chosen default plan. We use that RD design to estimate the effects of the change in default on the probability of switching plans and the probability of making an active choice. For those analyses, we study switching and active choice within 12 months of the change in default. We limit to only 12 months because, following that year, many plans that were benchmark plans in year t (and thus enroll both control beneficiaries and reassigned treated beneficiaries) may themselves lose benchmark status in year $t + 1$. Because of this, interpreting treatment effects in future years is tricky, as we must disentangle the treatment effect of the benchmark loss in t and the potential loss in $t + 1$.


In this section, we study the effects of the change in default on switching and active choice over 24 months following the implementation of the new default. To do so, we restrict the set of default plans we allow beneficiaries to be assigned to in year t to only those that do not lose benchmark status in $t + 1$. Specifically, for beneficiaries in plans losing benchmark status from year $t - 1$ to t , we restrict to those assigned to plans that retained benchmark status in $t + 1$. We implement this restriction based on the beneficiary's assigned plan rather than actual enrolled plan, as the assigned plan was randomly selected. Likewise, for the control group of beneficiaries in plans not losing benchmark status from year $t - 1$ to t , we restrict to those whose plan of enrollment as of December of year $t - 1$ retains benchmark status in both t and $t + 1$. After these restrictions, we are left with 304,602 beneficiary-experiment records in our restricted sample.

We replicate our difference-in-differences analyses with this new sample. Figures A15 and A16 present event study estimates of the effects of the change in default on drug spending, with Figure A15 presenting overall effects and Figure A16 presenting effects stratified by plan fit. These results indicate that the spending reductions from the change in the default persist through the end of the second year after the new default was implemented. Further, they show that the difference in the consequences of being assigned to a high-fitting versus a low-fitting plan are also persistent, implying that beneficiaries don't fully adapt to their new default plan's formulary or respond by switching to a plan with a smaller effect on drug spending.


Figures A17 and A18 present event study estimates of the effects of the change in default on the probability of switching plans and the probability of making an active choice, respectively. Again, the effects are persistent. Beneficiaries who faced a change in default switched out of their incumbent plans at extremely high rates and did not switch back within 24 months of the implementation of the new default. They also were not induced by the change in default to make an active choice. The rate of active choice was extremely low after the change in default and did not increase over time. Further, there is effectively no difference in effects on switching or active choice for beneficiaries assigned to higher- versus lower-fitting plans. Again, these results imply that beneficiaries did not respond to the negative effects of the new default by switching plans or making active choices later. Instead, they just passively followed the new defaults and absorbed the persistent reductions in drug spending.




Appendix Figure A15: This figure plots event-study regression estimates of the effect of reassignment to a new plan on quarterly drug spending, for beneficiaries assigned to plans of differing ‘fit.’ This is a replication of Figure 8a with our two-year follow-up sample.



Appendix Figure A16: This figure plots event-study regression estimates of the effect of reassignment to a new plan on quarterly drug spending, for beneficiaries assigned to plans of differing ‘fit.’ This is a replication of Figure 9a with our two-year follow-up sample.



Appendix Figure A17: This figure plots the propensity of beneficiaries whose incumbent plan lost benchmark status to switch from the plan they were enrolled in during December of year $t - 1$, stratified by the fit of the plan they were randomly assigned as a default for year t , relative to beneficiaries whose plan did not lose benchmark status. This is a replication of Figure 10a with our two-year follow-up sample.



Appendix Figure A18: This figure plots the propensity of beneficiaries whose incumbent plan lost benchmark status to make an active plan choice (not entering the default mechanism *or* switching plans after reassignment), stratified by the fit of the plan they were randomly assigned as a default for year t , relative to beneficiaries whose plan did not lose benchmark status. This is a replication of Figure 10b with our two-year follow-up sample.

D Appendix: Structural Exit Model Derivations

In Section 7, we would like to derive a closed-form expression for $P[a_{it}|a_{it'}, X]$. We do so in this section.

Recall that a depends on c and k . We have assumed by construct that k is independent across time periods, so $k_{it'}$ is irrelevant for a_{it} . Therefore, the value of $a_{it'}$ is the extent to which it is informative about c . We can start by explicitly characterizing this information by conditioning out c :

$$\begin{aligned} P[a_{it}|a_{it'}, X] &= \int_c P[a_{it}|c, a_{it'}, X] \cdot f(c|a_{it'}, X) dc \\ &= \int_c P[a_{it}|c] \cdot f(c|a_{it'}, X) dc \end{aligned}$$

where the second equality comes from the fact that, by conditioning on c , we have removed any additional information that $a_{it'}$ or X provides, since both are orthogonal to k_{it} , the remaining unexplained component of a_{it} . The second term is the information that $a_{it'}$ provides about the distribution of c given X . This is unknown, but we can decompose it using Bayes' rule:

$$f(c|a_{it'}, X) = \frac{P[a_{it'}|c, X] \cdot f(c|X)}{P[a_{it'}|X]}$$

and therefore, plugging this expression into the initial formula:

$$P[a_{it}|a_{it'}, X] = \frac{1}{P[a_{it'}|X]} \int_c P[a_{it}|c] \cdot P[a_{it'}|c] \cdot f(c|X) dc$$

The term before the integral is the share of beneficiaries with characteristics X who previously made an active choice, which can be computed directly from the data. The term in the integral has two unknowns. First, we need the conditional distribution of a given c . This requires us to make some assumption about the distribution of k . Second, we need the conditional distribution of c given X , which is an unknown that requires a parametric assumption.

In our empirical exercise, we make two distributional assumptions: That $c \sim \mathcal{N}(\mu X, \sigma^2)$, and that $k \sim \mathcal{N}(0, 1)$. The fixed mean and variance of the latter is required as a normalization, as is typical in parametric models of discrete choice. Given these assumptions, $P[a = 1|c] = P[c + k \geq 0] = 1 - \Phi(-c)$ and $P[a = 0|c] = P[c + k < 0] = \Phi(-c)$, while $f(c|X) = \phi\left(\frac{c - \mu X}{\sigma}\right)$.

With these assumptions, we can once again rewrite our conditional distribution, as

$$P[a_{it}|a_{it'}, X] = \frac{1}{P[a_{it'}|X]} \int_c [1 - \Phi(-c)]^{a_{it} + a_{it'}} \cdot [\Phi(-c)]^{2 - a_{it} - a_{it'}} \cdot \phi\left(\frac{c - \mu X}{\sigma}\right) dc$$

which is the closed-form expression for the likelihood of a_{it} given $a_{it'}$, which can be used in a maximum likelihood estimation procedure to estimate μ, σ^2 :

$$\hat{\mu}, \hat{\sigma}^2 = \arg \max_{\mu, \sigma^2} \sum_i \log \left[\int_c [1 - \Phi(-c)]^{a_{it} + a_{it'}} \cdot [\Phi(-c)]^{2 - a_{it} - a_{it'}} \cdot \phi \left(\frac{c - \mu X}{\sigma} \right) dc \right] - \log [P[a_{it'} | X]]$$

where the latter term does not depend on μ, σ^2 and therefore drops out.

Note that the integral within this expression has no closed form. We approximate it using Gauss-Hermite quadrature. This entails approximating the integral with

$$\frac{1}{\sqrt{\pi}} \sum_{k=1}^n w_k \left[1 - \Phi(-\mu - \sqrt{2}\sigma x_k) \right]^{a_{it} + a_{it'}} \cdot \left[\Phi(-\mu - \sqrt{2}\sigma x_k) \right]^{2 - a_{it} - a_{it'}}$$

where x_k is are the roots of the Hermite polynomial of degree n , $H_n(x)$, and

$$w_k = \frac{2^{n-1} n! \sqrt{\pi}}{n^2 [H_{n-1}(x_k)]^2}$$

Preliminary tests suggested that this approximation appears to converge at merely $n = 5$ nodes, for reasonable parameter values. For safety, we approximate using $n = 100$ nodes instead. We estimate standard errors in Table 7 using the bootstrap method.