

NBER WORKING PAPER SERIES

SOCIAL DEFAULTS AND PLAN CHOICE:
THE CASE OF SPOUSAL FOLLOWING

Tal Gross
Timothy Layton
Daniel Prinz
Julia Yates

Working Paper 34137
<http://www.nber.org/papers/w34137>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025

We are grateful to Keith Ericson, Jim Rebitzer, Ágnes Szabó-Morvai, and seminar participants at Lehigh University, Purdue University, Stanford University, the University of Georgia, Georgia State University, the University of Ottawa, the University of Pennsylvania, the University of Virginia, the University of San Diego, the University of Duisburg-Essen, and the CERS Institute of Economics for useful feedback. Amanda Stiebris provided superb research assistance. Research reported in this publication was supported by the National Institute on Aging of the National Institutes of Health under Award Number P30AG012810 through the National Bureau of Economic Research Center for Aging and Health Research. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not represent the views of the World Bank and its affiliated organizations, nor those of the Executive Directors of the World Bank, nor the governments they represent, nor the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34137>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Tal Gross, Timothy Layton, Daniel Prinz, and Julia Yates. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Social Defaults and Plan Choice: The Case of Spousal Following
Tal Gross, Timothy Layton, Daniel Prinz, and Julia Yates
NBER Working Paper No. 34137
August 2025
JEL No. D12, D14, I13

ABSTRACT

We study how couples in the Medicare Part D program choose an insurance plan. Over 70 percent of enrollees choose the same plan as their spouse. Even among those with differing healthcare needs, well over half do so. Discrete-choice models suggest beneficiaries value being on the same plan as their spouse at over a thousand dollars per year. Using a regression-discontinuity design, we show that younger spouses disproportionately follow their older spouse's plan choice. Joint plan choice contributes modestly to overall overspending, but increases costs substantially for the couples with different cost-minimizing plans.

Tal Gross
Boston University
Questrom School of Business
and NBER
talgross@bu.edu

Timothy Layton
University of Virginia
and NBER
timothyjlayton@virginia.edu

Daniel Prinz
The World Bank
dprinz@worldbank.org

Julia Yates
University of Michigan
jryates@umich.edu

Default options guide many important economic decisions. Many consumers, for instance, choose a retirement-savings plan passively by opting into whatever plan is the default option (Madrian and Shea, 2001; Beshears et al., 2008). Similarly, many people choose whether or not to become organ donors based on whether or not they are automatically enrolled in an organ-donation program (Johnson and Goldstein, 2003).

Other decisions require an active choice. Americans are not automatically assigned to a car insurer. They must choose one if they want to drive legally. In many contexts, the choice of a health insurance plan is similar: consumers are not automatically enrolled in a particular plan and so have to make an active choice among the plans offered to them.

Such active choices may facilitate better decisions (Carroll et al., 2009; Bernheim et al., 2015; Goldin and Reck, 2022). Without automatic enrollment in a default option, perhaps consumers are forced to weigh all options equally and so are more likely to choose the product that is best for them.

Then again, the absence of a default option does not imply that consumers make such decisions in a vacuum. They may use heuristics to make these decisions. In particular, consumers observe the choices of their peers, and their peers’ choices may act as a “social default.”¹ Instead of a social planner singling out a particular option as the default, the choices of peers may act in a similar fashion.

For complex products, social defaults may be problematic. It is one thing if consumers choose the same brand of packaged food as their neighbor (Bronnenberg et al., 2012) – it is quite another if they choose the same insurance plan. Insurance plans are often horizontally differentiated, and so the optimal plan for one person may be different than the optimal plan for another. In such settings, following a peer may lead a consumer to a product that is a sub-optimal fit for them, even if it is the right product for their peer.

In this paper, we study an active choice over a complex product in which most consumers are not initially offered a default option. Medicare recipients who join the Medicare Part D program must choose a private drug plan.² New Medicare recipients are presented with roughly three-dozen plans and must make an active choice between

¹The phrase “social default” was coined by Huh et al. (2014) after a series of laboratory experiments. The authors argue that “when a person is deciding between options for which her preferences are not well-formed, observing the choice of another person makes the option chosen by the other person a *social default*” (emphasis theirs).

²Medicare recipients who automatically qualify for the Low Income Subsidy program do face a default option (Ericson, 2014; Brot-Goldberg et al., 2023). We do not study those recipients here.

them.³ We study whether the choices of their peers affect their own choice of plan. To do so, we focus on a person’s closest peer, their spouse.

Our empirical analysis proceeds in three steps. First, we present a descriptive analysis to measure the share of couples on the same plan. Over 70 percent of couples in the Medicare Part D program are enrolled in the same plan. One driver of that behavior may be homophily: couples tend to have similar prescription-drug needs and preferences. And yet, even among couples with very different healthcare needs, we still find that a majority choose the same plan. Further, we identify “placebo spouses” via exact matching. Fewer than 20 percent of enrollees are on the same drug plan as their placebo spouse. Those patterns suggest that homophily – at least in terms of observable characteristics – is unlikely to explain the extremely high rate at which spouses choose the same plan.

Second, we estimate discrete-choice models. The models allow beneficiaries’ willingness-to-pay for each plan to vary according to their own expected out-of-pocket costs in that plan. We include in the model an indicator for whether the beneficiary’s spouse is enrolled in a given plan, allowing us to estimate how much the beneficiary is willing to pay to be on the same plan as their spouse. Those models suggest that younger spouses value being on the same plan as their spouse at about 76 percent of the value they place on remaining in the same plan from one year to the next (that is, inertia). We estimate that beneficiaries value keeping the previous year’s plan at \$2,294, in line with prior estimates that range from \$1,000 to \$2,000 per year (Handel, 2013; Polyakova, 2016). Our estimates suggest that beneficiaries are willing to forgo approximately \$1,741 per year in order to be on the same plan as their spouse.

The third step of our analysis takes a quasi-experimental approach involving a pair of natural experiments. Humana, a large insurer, twice introduced new plans to the market that immediately became very popular. Via a regression-discontinuity design, we find that those who joined Medicare Part D just before the new Humana plans were introduced were less likely to join those plans as compared to those who joined Medicare Part D just after the new Humana plans were introduced. That pattern is likely driven by inertia: incumbent Medicare recipients choose a plan that is available when they join the program and then stay in that plan despite the entrance of the new Humana plans.

We then study younger spouses who enter Medicare once the new Humana plans are available. Some younger spouses have older partners who entered Medicare just

³The precise number of plans varies by region and year.

before the new Humana plans were introduced – other younger spouses have older partners who entered Medicare just after the new Humana plans were introduced. Remarkably, we find that younger spouses are more likely to choose a Humana plan if their older partners joined Medicare just after the plans were introduced.

Those results cannot be explained as a consequence of spousal homophily. Instead, they imply that spouses simply want to be on the same plan. They may choose the same plan in response to the real costs associated with being on different plans or due to a behavioral bias. Ultimately, the evidence we compile cannot isolate the mechanism. That said, behavioral biases may be relevant. The discrete-choice models, in particular, suggest that spouses exhibit a willingness-to-pay to be on the same plan as their partners on the order of \$1,741. Such estimates seem likely to exceed any real hassle costs of being on different plans.

If spouses do indeed choose the same plan due to a behavioral friction, it is natural to wonder whether that friction may explain results from the prior literature on plan choice. Many studies document errors in Part D plan choice ([Abaluck and Gruber, 2011](#); [Ketcham et al., 2012](#); [Abaluck and Gruber, 2016](#); [Polyakova, 2016](#); [Ho et al., 2017](#); [Abaluck and Adams-Prassl, 2021](#); [Heiss et al., 2021](#); [Abaluck and Gruber, 2023](#); [Brot-Goldberg et al., 2023](#)). Are such errors driven by spouses choosing plans together? To investigate that possibility, we estimate the out-of-pocket costs that each beneficiary would have saved were they to have chosen a cost-minimizing plan. Previous researchers have performed similar exercises and have found that Part D beneficiaries could save approximately \$300 each year by choosing different Part D plans. We replicate that basic finding and then extend it with a novel decomposition. We decompose the total out-of-pocket costs that beneficiaries could save into two components. First, we estimate the amount of out-of-pocket costs that an individual might save by choosing their “household-cost-minimizing plan,” the plan that would lead to the lowest out-of-pocket costs conditional on the couple being on the same plan. Second, we estimate the additional money that the beneficiary would save by shifting from the household-cost-minimizing plan to the lowest-cost plan for them as an individual.

Most of the foregone savings in out-of-pocket costs are driven by the first component. On average, Part D beneficiaries could capture most of the possible savings by choosing a different plan *with their spouse*. This suggests that a large portion of foregone savings are not driven by “spousal following” per se, but rather by the couple jointly choosing a sub-optimal plan. That said, most beneficiaries have the

same individual- and household-cost-minimizing plans. Among beneficiaries who do not have the same individual- and household-cost-minimizing plans, substantial savings could be achieved by choosing separate plans. That group makes up only 38% of households, with the other groups lowering the average additional savings from optimal individual choice.

A final pattern in the discrete-choice models leads to a similar conclusion. Previous studies have used discrete-choice models to demonstrate that Part D enrollees are more sensitive to premiums than out-of-pocket costs. We first demonstrate the same result. We then show that the pattern is generally unchanged when we include in the model an indicator variable for one’s spouse being on that plan. That suggests that, at least on average, spousal following cannot explain much of the excess sensitivity to premiums over out-of-pocket costs.

This paper’s findings speak to the vast literature on insurance plan choice in general (e.g., [Sydnor, 2010](#); [Sinaiko and Hirth, 2011](#); [Handel, 2013](#); [Loewenstein et al., 2013](#); [Handel and Kolstad, 2015](#); [Bhargava et al., 2017](#); [Ericson and Sydnor, 2017](#)) and plan choice in the Medicare Part D program in particular ([Abaluck and Gruber, 2011](#); [Kling et al., 2012](#); [Ketcham et al., 2012](#); [Ericson, 2014](#); [Abaluck and Gruber, 2016](#); [Ho et al., 2017](#); [Abaluck and Adams-Prassl, 2021](#); [Heiss et al., 2021](#); [Abaluck and Gruber, 2023](#); [Brot-Goldberg et al., 2023](#)). We are aware of only a single, unpublished study that has documented the tendency of American couples to choose the same Medicare Part D plan ([Sanguinetti, 2019](#)).

1 Background and Data

We study a 20-percent sample of recipients enrolled in Medicare Part D, a federal program that provides insurance for prescription drugs to elderly and disabled Americans. Recipients choose among private plans, with roughly three-dozen options to choose from, depending on the state and year. Choice occurs at the time of Medicare entry and during an open-enrollment period toward the end of each year. Medicare entry often occurs in the month an individual turns 65, but some beneficiaries opt to enroll later. Critically, with the exception of those entering the Low Income Subsidy program, there is no default option at the time of initial enrollment.⁴

Which drug plan is best depends on the recipient. Plans differ in which drugs they cover, and so a first task for recipients choosing a plan is to estimate the total

⁴We exclude enrollees in the Low Income Subsidy program in what follows along with those on Employer Group Waiver Plans.

cost of each plan given the drugs that they have already been prescribed. Plans also differ in the insurance value that they offer based on the drugs that the recipient may be prescribed in the future. That insurance value depends on the idiosyncratic risk faced by the recipient along with the financial structure of the plan: whether it has a deductible, donut-hole coverage, etc. Finally, plans differ along other dimensions, including the pharmacies in their network and their brand name.

All of those factors make the choice complex. For that reason, the Centers for Medicare & Medicaid Services (CMS) offers a phone line and an online tool to help recipients choose a plan. Nevertheless, a large research literature has debated the degree to which recipients choose plans well.

Recipients choose plans amongst those available within their region. CMS operates 34 regions for Part D that cover all 50 states and the District of Columbia and regulates the plans on offer in each region. National and local insurance companies operate the plans. These insurance companies must offer a plan at or equivalent to a published standard coverage level (referred to as “basic” plans) and may offer plans with higher levels of coverage at typically higher costs (referred to as “enhanced” plans). For example, CVS-operated SilverScript offered the basic plan of “SilverScript Basic” and the enhanced plans “SilverScript Choice” and “SilverScript Plus” listed in Appendix Figure A1. Those multiple plans typically form tiers with varying premiums and levels of coverage.

We rely on two main datasets for our analysis: the Medicare Beneficiary Summary File and the Medicare Part D Event File. Those datasets cover the years 2006 through 2017, and we observe each recipient’s ZIP Code and their enrollment in Medicare Parts A, B, C, and D. We also use the Medicare Part D Formulary and Plan Characteristics files.

In order to identify spouses, we rely on a novel addition to these data. At our request, CMS provided us with encrypted identifiers of the address of residence for Medicare recipients enrolled in 2010 through 2017. We use those encrypted addresses to identify spouses. Appendix B describes that data.

2 Descriptive Analysis

2.1 Share of Couples on the Same Plan

The first column of Table 1 presents the share of couples on the same Part D plan across different sub-groups. Overall, over 72.6 percent of couples in the data are on

the same Part D plan.⁵ That number is larger for couples in which the older spouse is White than for couples in which the older spouse is non-White. The share of couples on the same plan is also roughly two percentage-points larger among those who reside in high-income ZIP Codes, suggesting that higher-income households are slightly more likely to choose together.⁶ When the older spouse’s drug spending is above the median, only 70% of couples pick the same plan, as compared to 75% when the older spouse’s spending is below the median.⁷ All of this suggests that the share of couples on the same plan is high across all groups, and even slightly higher among high-socioeconomic-status and healthier groups.

The remaining rows of Table 1 explore one particular dimension on which spouses may differ from each other. Among couples in which neither spouse has a chronic condition, 76 percent are on the same Part D plan. For couples in which both spouses have a chronic condition, 73.1 percent of couples are on the same Part D plan. Finally, there are couples in which precisely one spouse has a chronic condition: 69.4 percent of those couples are on the same Part D plan. Those three groups comprise a useful stratification. Couples in the last group are more likely to have different healthcare needs from one another and thus have different optimal plan choices. And yet, a similar share of couples across all three groups have chosen the same drug plan.⁸

In order to offer a baseline for comparison, the second column of Table 1 lists the percent of “placebo” couples choosing the same plan. To construct placebo spouses, we take all Medicare recipients in the data and match them to a nearest neighbor. Specifically, we perform nearest-neighbor matching based on: year, state, age, gender, White or non-White race, number of chronic conditions, and the ventile of median household income in their ZIP Code of residence. We then create placebo couples by linking each older Medicare recipient to their actual spouse’s nearest neighbor. The

⁵As an illustration of how couples sort across plans, consider Appendix Figure A1. That figure lists all major plans in Delaware in 2014, and demonstrates that couples tend to choose the same plan. The phenomenon is not limited to a single plan or segment of the market.

⁶Appendix Figure A2 plots the share of recipients on the same drug plan as their spouse by the ventile of the median household income in the recipients’ ZIP Code of residence. The figure suggests that choosing the same drug plan as one’s spouse is not a phenomenon exclusive to one income group.

⁷Appendix Figures A3 and A4 plot the same outcomes by deciles of drug spending and number of same chronic conditions.

⁸Appendix Figure A5 offers another stratification that leads to the same conclusion. That figure studies solely younger spouses joining Medicare because they have turned 65. Among those whose older spouse is less than five years older than them, over 55 percent choose the same plan as their older spouse. Among those whose older spouse is more than five years older than them, fewer choose the same plan.

percent of placebo spouses on the same plan represents the rate we would expect if plan choices were solely driven by observable characteristics. For all groups, actual couples choose the same plan at much higher rates than placebo couples. The share of actual couples on the same plan always exceeds 60 percent, while the rate for placebo couples is always below 20 percent.

To further explore that pattern, we study drug spending directly. Figure 1 stratifies couples by the difference in their annual drug spending, the older spouse’s spending minus the younger spouse’s spending, and plots the share of couples on the same Part D plan by ventiles of that difference. The solid circles suggest a modest inverted-V-shaped pattern. The couples with the most similar levels of drug spending have the highest likelihood of being on the same Part D plan. But even among couples with very different levels of drug spending in either direction, the share of couples on the same Part D plan is well above 60 percent. The hollow circles in Figure 1 plot the share of placebo couples on the same plan. The figure demonstrates that Medicare recipients are roughly three times less likely to be on the same plan as their placebo spouse compared to their actual spouse.

Overall, Table 1 and Figure 1 suggest that spousal following is pervasive and exists across various groups. In addition, the phenomenon is not one that is disappearing. Appendix Figure A6 plots the share of recipients on the same drug plan as their spouse by year that the recipient joined Medicare. Our data allows us to study spouses in 2010 through 2017, and the share of recipients on the same plan as their spouse is higher than 66 percent for recipients joining Medicare throughout those years.

2.2 Discrete-Choice Models

Previous studies of plan choice have relied on conditional logistic regressions in order to model consumers’ choice of plans. The regressions typically take the following form: the utility of plan j for Medicare recipient i is

$$u_{ij} = \beta_0 \cdot \pi_j + \beta_1 \cdot x_j + \beta_2 \cdot \text{OOP}_{ij} + \varepsilon_{ij}.$$

Here, π_j indicates plan j ’s premium, OOP_{ij} indicates person i ’s out-of-pocket expenses on plan j , and x_j indicates other characteristics of plan j , such as whether it features donut-hole coverage, a deductible, the plan’s brand, etc. We follow the literature in estimating such models, but restrict the sample to younger spouses and include covariates that account for their older spouses’ choices. We estimate out-of-

pocket spending for each individual in each available plan.⁹

Table 2 presents estimates of such regressions. The first column only includes controls for the premium, the presence of a deductible, and donut-hole coverage. Unsurprisingly, Medicare recipients are less likely to choose plans with higher premiums. Next, we include a control for out-of-pocket costs. The third column adds a key additional covariate: an indicator variable equal to one if the recipient’s older spouse has chosen that plan. The fourth column includes estimated out-of-pocket costs and the fifth column adds brand-specific fixed effects to the model. The addition of these variables increases confidence that the premium and “spouse on plan” variables are not capturing unobservable variation in quality or other demand-relevant factors across plans. The premium and “spouse on plan” coefficients are qualitatively similar across those specifications.

Finally, the sixth column of Table 2 includes a variable meant to capture recipients’ inertia: an indicator variable equal to one if the recipient was on the given plan during the previous year. As has been established by previous studies (Handel, 2013; Polyakova, 2016), the estimates suggest substantial inertia. Medicare recipients are much more likely to be on a plan if it was the plan they were on during the previous year. The measure of inertia provides another benchmark with which to judge the degree to which spouses choose the same plan. The “spouse on plan” coefficient is roughly 76 percent of the magnitude of the inertia coefficient. Prior work has shown that consumers leave substantial money on the table via inertia. These estimates thus suggest that the degree to which spouses choose the same plan is of roughly similar magnitude in its impact on recipients’ plan choice.

The ratio of the coefficients in that sixth model suggests that Medicare recipients are willing to pay $3.3409/0.001919 \approx 1,741$ dollars each year in premiums to be on a plan that their spouse has chosen. By way of comparison, average premiums were roughly \$30 each month, or \$360 annually, for basic plans in 2010 through 2016.

One pattern in Table 2 is worthy of note. Prior work, cataloged above, has identified that consumers place more weight on premiums relative to out-of-pocket costs. The second column of Table 2 replicates that result: the premium coefficient is indeed larger than the coefficient on out-of-pocket costs. And yet, when the model in the fourth column includes a control for “spouse on plan,” the gap between the

⁹To create the counter-factual out-of-pocket spending, we estimate the cost of every prescription an individual fills in every available plan based on the coverage phase at the time of the fill, the pharmacy type, the cost type (coinsurance versus copay), the drug tier of the filled prescription, and the days supply. This process is discussed in more depth in Appendix C.

premium and out-of-pocket-cost coefficients decreases by nearly 40%. That suggests that a portion of the pattern in the previous literature may be driven by the joint-plan-choice patterns this paper documents.

3 Quasi-Experimental Analysis

Consider an innovation in the Part D marketplace: the introduction of Walmart-branded plans. Humana, a large insurer, had operated Part D plans since the inception of the program in 2006. In 2011, Humana introduced a basic Humana-Walmart plan that involved both very low premiums (\$14.80 each month) and low copayments.¹⁰ As a result, the plan immediately proved popular. Then, in 2014, Humana renamed the 2011 plan the “Humana-Walmart Preferred RX Plan” and created a new, enhanced plan with an even-lower monthly premium: \$12.60 per month. That monthly premium was lower than any other enhanced plan’s premium – average premiums had been \$50–\$75 per month beforehand. That second, Humana-Walmart plan proved especially popular and increased the overall share of Medicare recipients on Humana plans.

But not all Medicare recipients were equally likely to join the new Humana plans. The first new Humana-Walmart plan was introduced in January of 2011 and the second new Humana-Walmart plan was introduced in January of 2014. Older spouses who joined Medicare in early 2011 or 2014 were more likely to join the new Humana plan than older spouses who joined Medicare in late 2010 or 2013. That contrast is likely driven by inertia. Entrants in 2011 and 2014 had to make an active choice across all plans and so were likely to choose the new Humana plan. By contrast, entrants in 2010 and 2013 were not presented with the new Humana plan when they first entered Medicare and did not have to make an active choice the following year. To join the new Humana plan, they would have had to choose some other plan in their year of entry and then decide to actively switch to the new Humana plan the following year instead of passively re-enrolling in their 2010 or 2013 plan. As a result, an older spouse joining Medicare in early 2011 or 2014 was more likely to be on Humana than an otherwise-similar older spouse joining Medicare in late 2010 or 2013.

Appendix Figure A8 describes this variation. The first panel plots the share of new entrants onto Medicare who chose a Humana plan each year. The graph demonstrates a discrete increase in the share choosing Humana plans in 2011, the year that Humana

¹⁰As an illustration, Figure A7 plots mean premiums for plans in the Texas market.

introduced its first lowest-price Humana-Walmart plan, and then again in 2014, the year that the second Humana-Walmart plan was introduced. The second panel plots the share of younger spouses who entered Medicare in 2016 separately by the year that their older spouse joined Medicare. The figure suggests similar, sharp increases in the likelihood of choosing Humana at Medicare entry in 2016 for younger spouses whose older spouse entered Medicare in 2011 versus 2010 and again in 2014 versus 2013.

This setting lends itself to a regression-discontinuity design. Consider older spouses who join Medicare just before or after a new Humana plan was introduced in 2011 and 2014. The running variable is the difference between the month that all older spouses turn 65 and the January that a new Humana-Walmart plan was introduced. We then test for a discontinuity in Humana enrollment first for the older spouse and then for their younger partner.

Older spouses with 65th birthdays just before the first of January of 2011 or 2014 are likely similar to older spouses with 65th birthdays just after those dates. More importantly, the younger spouses of those older spouses are likely similar as well. In that sense, such an analysis eliminates the possibility that differences in plan choices are driven by underlying differences in the preferences of the older or younger spouses.

There exists, however, a complication. The new Humana-Walmart plans were introduced in January – all new plans are introduced at the start of the year. The demographics of Medicare recipients change discontinuously at the start of the year. For instance, we observe that Medicare recipients who are born in January are less likely to be White. (Appendix Figures [A9](#) and [A10](#) present those patterns.) Those demographic differences may be driven by seasonal fertility patterns. They may also be driven by Medicare recipients who do not know their actual birthday and so are coded as having a birth date in January.

Given that complication, we implement a “regression-discontinuity difference-in-difference” (RD-DD) approach inspired by the work of [Persson and Rossin-Slater \(2024\)](#). The specification on which we rely compares older spouses who turned 65 just before January 2011 or 2014 to other older spouses who turned 65 just after January 2011 or 2014, and then makes that same comparison for other years. That approach differences out the annual seasonal variation around the start of the year. For all older spouses i , whose 65th birthday falls d_i days after January first of year t ,

we estimate:

$$\begin{aligned} \text{Humana}_i = & \alpha_0 + \alpha_{t(i)} + \beta \cdot \mathbf{1}[d_i \geq 0] \cdot \mathbf{1}[t = 2011] + \\ & \gamma_1 \cdot \mathbf{1}[d_i \geq 0] + \gamma_2 \cdot f(d_i) + \gamma_3 \cdot \mathbf{1}[d_i \geq 0] \cdot f(d_i) + \varepsilon_i, \end{aligned}$$

where Humana_i is an indicator equal to one if older spouse i born d days after January first in year t is observed to be enrolled in a Humana plan. (The regression equation for the 2014 change is identical, though with the key interaction term of interest being $\beta \cdot \mathbf{1}[d_i \geq 0] \cdot \mathbf{1}[t = 2014]$.) We then run an identical specification but with an outcome that measures the Humana enrollment of person i 's younger spouse. This specification relies on a quadratic control for the running variable and a bandwidth of 90 days before and after January first.¹¹ While a typical regression-discontinuity specification relies on the assumption that those on either side of the discontinuity are identical, this RD-DD specification relies on the weaker assumption that any differences across the discontinuity are systematic and persistent across years.

Appendix Figure A11 presents graphical estimates of the underlying RD specifications for older spouses' enrollment on Humana. The first and third panels present the main estimates for 2011 and 2014 respectively. Older spouses are more likely to be enrolled on Humana if their birthdays fall just after the first of January in 2011 or 2014 as compared to just before the first of January in those years. The second and fourth panels present the same RD but for years other than 2011 and 2014. Those figures suggest a discontinuity in placebo years, but one that is smaller in magnitude and of the opposite sign as the main effects. Appendix Figure A12 presents graphical estimates of the underlying RD specifications for *younger* spouses' enrollment on Humana plans. The figures suggest that younger Medicare recipients whose older spouses turned 65 just after the introduction of Humana-Walmart plans were more likely to be on Humana themselves.

The top two panels of Figure 2 present RD-DD estimates of what might be called the first stage of this case study. The outcome of interest for those figures is the share of older spouses who choose a Humana plan relative to the month that they turned 65 minus the nearest January. Each gray circle in those panels plots the difference between the mean of that outcome for that running variable in 2011 or 2014 and the mean in other years. Both panels suggest a discontinuity in January. Older spouses who turned 65 just after the new Humana plans were introduced in 2011 and 2014

¹¹Appendix Tables A1 and A2 presents estimates of this specification with different bandwidths. The main estimates change little across varying bandwidths.

were more likely to choose a Humana plan.

The bottom two panels of Figure 2 present RD-DD estimates of what might be called the reduced form of this case study. The outcome of interest for those panels is the share of *younger* spouses who chose Humana based on the 65th birthday of their older partner. Younger spouses are more likely to be on Humana if their older spouse had a 65th birthday just after the introduction of a Humana-Walmart plan, than if their older spouse had a 65th birthday just before the introduction of a plan. The discontinuity is roughly 7 percentage points, which is roughly one third of the pre-existing average Humana enrollment rate.

Appendix Figure A13 plots those main RD-DD estimates against RD-DD estimates for pre-existing characteristics of younger spouses. As placebo tests, we estimate the effect of the RD-DD on race, income, and gender. Reassuringly, all RD-DD estimates for those pre-existing characteristics are statistically indistinguishable from zero. This supports the identifying assumption that differencing the RD across years allows us to adjust for demographic changes that occur in January.

4 Overspending Due to Joint Plan Choice

This section presents a decomposition to assess how much joint plan choice accounts for the overspending that has been documented in the previous literature. First, we determine the plan that would minimize ex-post out-of-pocket costs for each beneficiary. We then estimate the spending that each beneficiary would have saved if they had chosen this personal optimum plan. The first row and first column of Table 3 presents that number. The estimate, \$690, is larger than estimates in the previous literature, though of the same magnitude.¹²

Second, we calculate for each couple the plan that would minimize ex-post out-of-pocket costs for the couple as a whole, so long as they choose a plan together. That is an exercise in constrained optimization: choosing the cost-minimizing plan subject to the constraint that spouses remain on the same plan. The second row and first column of Table 3 presents the estimated savings if the recipient had chosen their joint optimum plan. On average, couples on the same plan could save \$614, ex post, were they to choose a different plan together.

Those two numbers imply that, on average, married Medicare recipients could save only an additional 12 percent in out-of-pocket costs if they chose a different plan

¹²The sample in Table 3 is limited to those who choose the same plan as their spouse, and that may drive any difference with previous estimates.

than their partner rather than choosing the plan that is best for the couple together. This, however, conceals important heterogeneity.

The sample consists of four groups. First, many recipients have the same personal optimum plan as their spouse. For that group, each personal optimal plan is also the joint optimal plan. The second column of Table 3 shows the decomposition for those couples. Unsurprisingly, the two overspending numbers are identical, because, by definition, such couples gain nothing from choosing a different plan from their spouse.

At the other extreme are couples for whom neither partner’s personal optimum matches the joint optimum. For them, picking one common plan forces both spouses off their optimum. The third column of Table 3 performs the same exercise for those couples. For such couples, the estimates suggest much more, an additional 36 percent, to be gained by choosing separate plans.

There are two remaining groups of recipients. We refer to one group as “dominant spouses:” those whose personal optimum coincides with the joint optimum. Their partners do not share that optimum and are therefore “dominated” spouses. The fourth column of Table 3 presents the decomposition for dominant spouses and the final column presents the decomposition for their partners, the “dominated” spouses. Those final numbers suggest that among “dominated spouses,” an additional 87 percent of savings can be achieved by choosing separate plans.

Table 3 therefore illustrates two implications of joint plan choice. First, couples’ tendency to choose the same plan explains a small share of *overall* overspending. Almost all of the possible savings in out-of-pocket costs can be achieved by couples choosing a cost-minimizing plan jointly. But the conclusion is very different for the 38-percent of the sample that has a different optimum plan as their spouse. For that group, joint plan choice substantially increases the out-of-pocket costs they could otherwise avoid. In that fashion, these results suggest substantial savings may be possible in settings in which peers are more likely to have different optimal plans.

5 Discussion

We find that Medicare enrollees tend to choose the same drug plan as their spouse. That finding is not solely a consequence of homophily. A remarkably high share of couples are on the same plan even among those with differing healthcare needs. Quasi-experimental evidence further suggests that spousal following is not a consequence of

unobservable characteristics of couples or plans. Younger spouses whose older spouses were born just weeks apart make different choices based on whether a given plan was available at the time their older partners turned 65.

Couples may believe that choosing the same plan lowers their administrative burden. Part D plans, however, require relatively little administration on the part of the recipient. In addition, discrete-choice models suggest that couples are willing to forgo strikingly large sums of money each month for such a hypothetical gain in reduced administration. Those sums are so large that they are difficult to rationalize.

A more likely mechanism, albeit one that we cannot isolate, involves cognitive overload. Many Medicare recipients likely dedicate relatively little bandwidth to their choice of plan. As such, when a plan proves to be satisfactory for an older spouse, younger spouses simply join it as well, rather than perform a full evaluation of the alternatives. It is also possible that spouses specialize in financial decisions, and only one spouse chooses plans, regardless of whom is gaining Medicare coverage. In either case, the key issue involves limited bandwidth. Such a phenomenon would be consistent with much of the broader literature on imperfect plan choice in the Part D program.

Finally, one can view the results of this paper through the lens of “peer effects:” how a person’s peers affect their consumption. [Bronnenberg et al. \(2012\)](#) and [Bailey et al. \(2022\)](#) demonstrate that consumers tend to choose the same brand or product as their peers. This paper presents a similar finding: consumers tend to choose the same insurance plan as their spouse. But, in contrast to those previous studies, the peer effect in this paper may be problematic. When a product’s quality is individual-specific, peer effects may lead many to make the wrong choice. For health insurance, the best product for my peer, may not be the best product for me.

References

- Abaluck, Jason and Abi Adams-Prassl**, “What do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses,” *Quarterly Journal of Economics*, 2021, *136* (3), 1611–1663.
- **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.
- **and –**, “Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance,” *American Economic Review*, 2016, *106* (8), 2145–2184.
- **and –**, “When Less is More: Improving Choices by Limiting Choices in Health Insurance Markets,” *Review of Economic Studies*, 2023, *90* (3), 1011–1040.
- Bailey, Michael, Drew M. Johnston, Theresa Kuchler, Johannes Stroebe, and Arlene Wong**, “Peer Effects in Product Adoption,” *American Economic Journal: Applied Economics*, 2022, *14* (3), 488–526.
- Bernheim, B. Douglas, 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 Choi, David Laibson, Brigitte Madrian, Stephen J. Kay, and Tapen Sinha**, “The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States,” in “Lessons from Pension Reform in the Americas,” Oxford University Press, 2008, pp. 59–87.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor**, “Choose to Lose: Health Plan Choices from a Menu with Dominated Option,” *Quarterly Journal of Economics*, 2017, *132* (3), 1319–1372.
- Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow**, “The Evolution of Brand Preferences: Evidence from Consumer Migration,” *American Economic Review*, 2012, *102* (6), 2472–2508.
- Brot-Goldberg, Zarek C., Timothy Layton, Boris Vabson, and Adelina Yanyue Wang**, “The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D,” *American Economic Review*, 2023, *113* (10), 2718–2758.
- 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.

- Ericson, Keith M. Marzilli**, “Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange,” *American Economic Journal: Economic Policy*, 2014, 6 (1), 38–64.
- Ericson, Keith Marzilli and Justin Sydnor**, “The Questionable Value of Having a Choice of Levels of Health Insurance Coverage,” *Journal of Economic Perspectives*, 2017, 31 (4), 51–72.
- Goldin, Jacob and Daniel Reck**, “Optimal Defaults with Normative Ambiguity,” *Review of Economics and Statistics*, 2022, 104 (1), 17–33.
- Handel, Benjamin R.**, “Adverse Selection and Inertia in Health Insurance Markets,” *American Economic Review*, 2013, 103 (7), 2643–82.
- **and Jonathan T. Kolstad**, “Health Insurance for ”Humans”: Information Frictions, Plan Choice, and Consumer Welfare,” *American Economic Review*, 2015, 105 (8), 2449–2500.
- Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou**, “Inattention and Switching Costs as Sources of Inertia in Medicare Part D,” *American Economic Review*, 2021, 111 (9), 2737–2781.
- 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.
- Huh, Young Eun, Joachim Vosgerau, and Carey K. Morewedge**, “Social Defaults: Observed Choices Become Choice Defaults,” *Journal of Consumer Research*, 2014, 41 (3), 746–760.
- Johnson, Eric J. and Daniel Goldstein**, “Do Defaults Save Lives?,” *Science*, 2003, 302 (5649), 1338–1339.
- Ketcham, Jonathan D., Claudio Lucarelli, Eugenio J. Miravete, and M. Christopher Roebuck**, “Sinking, Swimming, or Learning to Swim in Medicare Part D,” *American Economic Review*, 2012, 105 (2), 2639–73.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel**, “Comparison Friction: Experimental Evidence from Medicare Drug Plans,” *Quarterly Journal of Economics*, 2012, 127 (1), 199–235.
- Loewenstein, George, Joelle Y. Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, James J. Choi, Jonathan Kolstad, David Laibson, Brigitte C. Madrian, John A. List, and Kevin G. Volpp**, “Consumers’ Misunderstanding of Health Insurance,” *Journal of Health Economics*, 2013, 32 (5), 850–862.


- 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.
- Persson, Petra and Maya Rossin-Slater**, “When Dad Can Stay Home: Fathers’ Workplace Flexibility and Maternal Health,” *American Economic Journal: Applied Economics*, 2024, *16* (4), 186–219.
- Polyakova, Maria**, “Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D,” *American Economic Journal: Applied Economics*, 2016, *8* (3), 165–95.
- Sanguinetti, Tomas**, “How Do Couples Choose Individual Insurance Plans? Evidence from Medicare Part D,” Technical Report, Arizona State University Department of Economics 2019.
- Sinaiko, Anna D. and Richard A. Hirth**, “Consumers, health insurance and dominated choices,” *Journal of Health Economics*, 2011, *30* (2), 450–457.
- Sydnor, Justin**, “(Over)insuring Modest Risks,” *American Economic Journal: Applied Economics*, 2010, *2*, 177–199.

Table 1. Share of Part D Recipients on the Same Plan as Their Spouse

	Percent on Same Plan, Actual Couples	Percent on Same Plan, Placebo Couples
All recipients	72.6	17.5
White recipients	72.9	17.5
Non-White recipients	69.0	17.5
Low-income ZIP Code	71.1	16.3
High-income ZIP Code	73.3	18.1
Below-median drug expenditure	75.0	17.8
Above-median drug expenditure	69.7	17.1
No chronic conditions	76.0	19.0
Both spouses have conditions	73.1	17.5
One spouse has a condition	69.4	17.5

This table presents the share of Medicare Part D recipients who are on the same Part D drug plan as their spouse by sub-group. We measure income as the median household income of the ZIP Code of residence based on the American Community Survey. Race, income, and drug spending categories are defined based on the characteristics of the older spouse. Placebo couples are defined as the younger spouse matched to their older spouse's nearest neighbor based on observable characteristics.

Figure 1. Share of Couples on Same Plan by Difference in their Drug Spending



This figure plots the share of couples on the same Part D drug plan by the difference between the older spouse's annual drug spending and the younger spouse's annual drug spending. The solid circles plot the share of couples on the same drug plan for actual couples. The hollow circles plot the share of couples on the same drug plan when we match Medicare recipients to a placebo spouse based on exact matching as described in the main text. The lowest and highest ventile of spending are excluded for clarity.


Table 2. Models of Plan Choice

	(1)	(2)	(3)	(4)	(5)	(6)
Spouse on Plan			4.1036*** (0.0105)	4.0898*** (0.0106)	3.5616*** (0.0117)	3.3409*** (0.0108)
Premium (in hundreds)	-0.1561*** (0.0016)	-0.1660*** (0.0017)	-0.1167*** (0.0020)	-0.1261*** (0.0021)	-0.1574*** (0.0026)	-0.1919*** (0.0030)
OOP Cost (in hundreds)		-0.0683*** (0.0020)		-0.0664*** (0.0036)	-0.0599*** (0.0037)	-0.0617*** (0.0048)
On Plan Last Year						4.4016*** (0.0201)
Deductible	✓	✓	✓	✓	✓	✓
Donut Hole Coverage	✓	✓	✓	✓	✓	✓
Brand FE					✓	✓


This table presents conditional logistic regressions to assess plan choice. For the sake of computational time, these regressions are based solely on Part D recipients in California, Florida, New York, Pennsylvania, and Texas. There are 5,413,778 observations in the sample. The symbol * indicates that $p < 0.10$; ** indicates that $p < 0.05$; *** indicates that $p < 0.01$. We use the abbreviation “FE” for “fixed effect,” and “OOP Cost” for “out-of-pocket cost.”

Figure 2. Case Study of Humana Enrollment: RD-DD Estimates


(a) Older Spouses, 2011




(b) Older Spouses, 2014



(c) Younger Spouses, 2011



(d) Younger Spouses, 2014



These figures describe enrollment in Humana plans before and after the 2014 introduction of the second Humana-Walmart plan. All panels presents the difference-in-discontinuities design. The running variable is the difference between the older spouse's 65th birthday and January 1 of the relevant period. The first two panels show the difference in the share of older spouses who joined a Humana plan during the main treatment period (either cutoff at January 1, 2011 or January 1, 2014) and the share of younger spouses who joined a Humana plan in control periods with cutoffs on January 1st in the years pre and post the plan change. The second two panels show the difference in the share of younger spouses who joined a Humana plan during the main treatment period (either cutoff at January 1, 2011 or January 1, 2014) and the share of younger spouses who joined a Humana plan in control periods with cutoffs on January 1st in the years pre and post the plan change.


Table 3. Overspending Decomposition

	(1)	(2)	(3)	(4)	(5)
	Full	Same	Different	Dominant	Dominated
	Sample	Plan	Plan	Spouses	Spouses
Over-spending relative to <i>personal</i> cost-minimizing plan	690.11	412.06	633.78	1309.36	463.19
Over-spending relative to <i>household</i> cost-minimizing plan	614.13	412.06	467.47	1309.36	247.85
Share of sample	1.00	0.35	0.11	0.27	0.27

This table presents the average and median difference in out-of-pocket spending based on Part-D plan choice. The first row presents the average and median difference between each individual's out-of-pocket spending in their chosen plan and the plan in which they would have incurred the smallest out-of-pocket costs. The second row presents the average and median difference between each individual's out-of-pocket spending in their chosen plan and the plan where, if chosen by them and their spouse, they would have incurred the smallest out-of-pocket costs as a household. The first two columns present statistics for the entire sample and the second two columns present statistics solely for couples who do not have the same cost-minimizing plan.


A Appendix Tables and Figures

Appendix Figure A1. Example of Distribution of Couples in Delaware, 2014




As a case study, this figure lists all major Part D plans operating in Delaware in 2014. The figure lists the plans of older spouses (vertical list) and younger spouses (horizontal list). The coloring of each cell describes the share of couples in each combination of plans, with the shares adding to one across each horizontal line.

Appendix Figure A2. Spousal Following by ZIP Code Income




This figure plots the share of couples on the same Part D plan by the median household income in the couple's ZIP Code of residence.

Appendix Figure A3. Spousal Following Separately by Spending on Drugs




This figure plots the share of Part D recipients on the same drug plan as their spouse by decile of annual drug spending of the older spouse.

Appendix Figure A4. Spousal Following Separately by Similar Chronic Conditions




This figure plots the share of Medicare recipients who choose the same plan as their spouse based on how many chronic conditions the spouses share.

Appendix Figure A5. Spousal Following Separately by Age Difference




This figure plots the share of new Medicare entrants who choose the same plan as their older spouse, separately by the age of that older spouse.

Appendix Figure A6. Spousal Following Separately by Year




This figure plots the share of new Medicare entrants who choose the same plan as their older spouse, separately by the year in which they join Medicare.

Appendix Figure A7. Premiums in Texas Around the Humana-Walmart Entry



This figure plots mean monthly premiums for all Part D plans in Texas. In 2010, Humana offered two basic plans and one enhanced plan. In all other years, Humana offered one basic plan and two enhanced plans. The first Humana-Walmart plan was launched as a basic plan in January of 2011. The second Humana-Walmart plan was launched as an Enhanced plan in January of 2014.


Appendix Figure A8. Humana Enrollment: Annual Trends




These figures plot annual means for the Humana case study. The first panel presents the share of new Medicare enrollees who choose a Humana plan by the year they enter the Medicare program. The second panel plots the share of 2016 entrants to Medicare who choose a Humana plan based on the year that their older partner entered Medicare.

Appendix Figure A9. Humana Enrollment 2011 Case Study: Controls


(a) Share White, Treatment Period




(b) Share White, Control Period




(c) Share Female, Treatment Period




(d) Share Female, Control Period



(e) Share Above Median Income Zip Code, Treatment Period




(f) Share Above Median Income Zip Code, Control Period




This figure shows the regression discontinuity plots for the characteristics of younger spouses. The treatment period runs from October 2010 to March 2011. The control period includes October 2008-March 2009, October 2009-March 2010, October 2011-March 2012 and October 2012-March 2013.

Appendix Figure A10. Humana Enrollment 2014 Case Study: Controls




This figure shows the regression discontinuity plots for the characteristics of younger spouses. The treatment period runs from October 2013 to March 2014. The control period includes October 2011-March 2012, October 2012-March 2013, October 2014-March 2015 and October 2015-March 2016.

Appendix Figure A11. Regression-Discontinuity Estimates of the Effect of the Introduction of a New Humana Plan on Older Spouses' Enrollment in Humana




These figures plot the share of older spouses enrolled on a Humana plan on the vertical axis against their birthday relative to January first. The first two panels study the 2011 introduction of the basic Humana-Walmart plan and the second two panels study the 2014 introduction of the enhanced Humana-Walmart plan. The placebo tests for the 2011 Humana-Walmart introduction include older spouses whose birthdays were within a few months of January 2009, 2010, 2012, and 2013. The placebo tests for the 2014 Humana-Walmart introduction include older spouses whose birthdays were within a few months of January 2012, 2013, 2015, and 2016. The gray bars plot 95-percent confidence intervals for the means plotted.

Appendix Figure A12. Regression-Discontinuity Estimates of the Effect of the Introduction of a New Humana Plan on Younger Spouses' Enrollment in Humana



These figures plot the share of younger spouses on a Humana plan on the vertical axis against the relative birthday of their older spouses. The running variable is the 65th birthday of the older spouses relative to the first of January. The first two panels study the 2011 introduction of the basic Humana-Walmart plan and the second two panels study the 2014 introduction of the enhanced Humana-Walmart plan. The placebo tests for the 2011 Humana-Walmart introduction include older spouses whose birthdays were within a few months of January 2009, 2010, 2012, and 2013. The placebo tests for the 2014 Humana-Walmart introduction include older spouses whose birthdays were within a few months of January 2012, 2013, 2015, and 2016. The gray bars plot 95-percent confidence intervals for the means plotted.

Appendix Figure A13. Humana Regression-Discontinuity Difference-in-Difference Estimates



These figures plot regression-discontinuity difference-in-difference estimates for the 2011 (first panel) and 2014 (second panel) introduction of the Humana-Walmart plans. The outcome of each regression is listed in each panel: whether the younger spouse was observed on a Humana plan, whether the younger spouse self reported White race, and so on. “High income” indicates whether the ZIP Code of residence has a median household income that is above the overall median household income. The running variable for each regression is the 65th birth date of the older spouse. The coefficients plotted are interactions between an indicator variable equal to one if the older spouse’s birthday was after January first and an indicator variable equal to one if the January first in question was the year that the new plan was introduced (2011 or 2014). We define the treatment period as October through March surrounding January of 2011 or 2014. The control periods include those months surrounding January 2012, 2013, 2015, and 2016.

Appendix Table A1. Bandwidth Selection: Humana Enrollment 2011 Case Study

Bandwidth	30 Days	60 Days	90 Days	120 Days	150 Days
Younger spouse on Humana	0.0539*** (0.0103)	0.0548*** (0.0091)	0.0528*** (0.0085)	0.0516*** (0.0077)	0.0500*** (0.0069)
<i>N</i>	22,174	45,967	72,660	99,103	126,678

This table presents alternative bandwidths for the regression discontinuity difference-in-difference model. The symbol * indicates that $p < 0.10$; ** indicates that $p < 0.05$; *** indicates that $p < 0.01$.

Appendix Table A2. Bandwidth Selection: Humana Enrollment 2014 Case Study

Bandwidth	30 Days	60 Days	90 Days	120 Days	150 Days
Younger spouse on Humana	0.0617*** (0.0095)	0.0622*** (0.0062)	0.0595*** (0.0059)	0.0575*** (0.0056)	0.0556*** (0.0053)
<i>N</i>	27,006	56,077	88,832	120,476	153,166

This table presents alternative bandwidths for the regression discontinuity difference-in-difference model. The symbol * indicates that $p < 0.10$; ** indicates that $p < 0.05$; *** indicates that $p < 0.01$.

B Appendix: The Identification of Spouses


As described in the main text, we requested encrypted identifiers of every Medicare recipient’s address of residence. Analysts at the Centers for Medicare & Medicaid Services created those encrypted identifiers for us, with one encrypted identifier per beneficiary per year, from 2010 through 2017.

We process those encrypted identifiers via the following procedure. First, we eliminate individuals who appear at multiple addresses during the 2010–2017 period. Second, we eliminate from the data any encrypted addresses linked to more than two residents at any point in time. Those addresses are likely group homes, data errors, or unusual situations in which multiple households share the same mailing address. In addition, this restriction excludes instances where individuals remain at their address but change partners. We then categorize two Medicare recipients residing in the same address as spouses. We also restrict the sample to couples that both appear in the Medicare Part D 20-percent sample and neither are on a Medicare Advantage plan.

The identification of spouses in the data is an inference. Unfortunately, we do not observe marriage certificates or the self-reported identification of spouses. There may be some couples in the Medicare data that are not actually spouses but live in the same location. That said, the object of study here – joint plan choice – is still an interesting phenomenon among those households.

Reassuringly, though, the age differences among spouses we identify in the Medicare data matches a comparison group. Appendix Figure B1 studies households surveyed by the American Community Survey (ACS) in which the household contains two self-reported spouses who are both older than age 65. The difference in ages between spouses in the ACS and the difference in ages between spouses who are both Medicare recipients is strikingly similar.

Appendix Figure B1. Validation of Spousal Crosswalk based on Age



This figure plots percentiles of the age distribution across spouses measured in two samples. The blue circles plot the distribution among couples over the age of 65 that we identify in 1-year samples of the American Community Survey, the ACS, from 2013-2015. The red diamonds plot the distribution among couples over the age of 65 that we identify in the Medicare data from 2013-2015.

C Appendix: Estimation of Expected Out-of-Pocket Costs

In order to calculate the expected out-of-pocket costs for each spouse in each plan, we rely on the methodology and code published by [Abaluck and Gruber \(2011\)](#). We begin by identifying a sample of relevant couples from the Master Beneficiary Summary File. We start with the couples that we can identify by encrypted address and limit to those who are also in the Medicare Part-D 20-percent sample. To more-accurately predict the out-of-pocket costs faced by those couples, we exclude beneficiaries who ever received premium or cost-share subsidies either through a third party or the Low Income Subsidy program. For ease of computation, we also limit the sample to the five most-populated states: California, New York, Texas, Florida, and Pennsylvania. We further limit the sample to younger spouses who enter Medicare in a different year than their older spouse. Finally, we limit the sample to younger spouses who are in the same plan for all 12 months.

We merge the Medicare Part D Event file to enrollment data for that sample. The Part D Event file is comprised of every drug transaction covered by Medicare Part D. For each beneficiary’s transaction, the data identify the drug by National Drug Code (NDC), date of purchase, total and out-of-pocket cost, and the days supply.

We next construct information on plan formularies. The Medicare data include formulary files that provide information on the selection of drugs used by the full set of Part-D plans in a given year. Using these data, the Medicare Part D claims data, and other Medicare plan information, we create two crosswalks. The first crosswalk lists formulary information by NDC. For every NDC, we identify whether it falls into the formulary of a given plan. If it is included in the formulary, we identify which drug tier it falls under. Typically, a drug in a lower tier will cost a recipient less than a drug in a higher tier. The second crosswalk lists formulary and tier information by plan.

We then create a dataset detailing cost sharing. We start with the Medicare Part-D tier data. These data identify the cost-sharing rules by plan, tier, days supply and coverage phase. Note that Medicare Part D beneficiaries incur different cost-share amounts depending on their previous spending. During our sample period, there were typically four phases: deductible, pre-initial coverage limit or gap, gap, and catastrophic coverage. Using the tier data and the second formulary crosswalk, we identify the beneficiary cost-sharing information by plan, formulary, tier, days supply, coverage phase, and whether the drug is filled at a preferred pharmacy. Cost-sharing information may be listed as a copay amount or as coinsurance. For example, a drug

with 30-days supply in tier 1 filled during the pre-initial coverage limit may have a coinsurance amount of 25 percent of price or a copayment amount of \$4.

Using plan information, we next identify every plan available in each state and year. Merging these data with the sample, we create a dataset comprised of observations for each individual-NDC- prescription-fill-plan combination. To merge the cost-sharing crosswalk to these data, we merge on the first formulary crosswalk to identify the formulary and tier for each NDC-plan combination. We then merge on the cost-sharing data by tier, plan, days supply, and formulary. Finally, for each beneficiary, we sort their prescription fills by fill date and identify, for each plan, the cost-sharing amount for each fill based on days supply, tier, and coverage phase. We assume that each beneficiary would keep their prescription-fill history constant regardless of plan choice.