

Social Defaults and Plan Choice: The Case of Spousal Following*

Tal Gross[†] Tim Layton[‡] Daniel Prinz[§] Julia Yates[¶]

January 2026

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 hundreds of dollars per year. Using a regression-discontinuity design, we show that younger spouses disproportionately follow their older spouse's plan choice. We then present a decomposition exercise to assess how much the tendency for couples to choose the same plan contributes to overspending in the Part D program.

*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, or those of the Executive Directors of the World Bank or the governments they represent.

[†]Boston University and NBER, talgross@bu.edu

[‡]University of Virginia and NBER, timothyjlayton@virginia.edu

[§]World Bank, dprinz@worldbank.org

[¶]University of Michigan, jryates@umich.edu

It is difficult to choose a health insurance plan. Plans have different deductibles, copayments, and coverage. Many consumers struggle with such complex products, leading to costly mistakes in 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](#)). And yet, consumers do not make these choices in isolation: they observe the decisions of those around them. Consumers’ peers’ choices may serve as a “social default,” guiding their behavior.¹

In this paper, we study how consumers choose among drug insurance plans in the Medicare Part D program.² New Medicare recipients are presented with roughly three-dozen plans and must make an active choice between 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 77 percent of

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 – all other recipients must make an active choice with no plan singled out by Medicare as a default.

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

the value they place on remaining in the same plan from one year to the next (that is, inertia). Past research has estimated consumers' valuation of inertia at over a thousand dollars per year, suggesting that consumers highly value being 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 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 choose those plans relative 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 before the new Humana plans were introduced – other younger spouses have older partners who entered Medicare just after the new Humana plans were introduced. 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 and were thus themselves more likely to be in a Humana plan.

These results cannot be explained by spousal homophily. While homophily might lead some couples to both choose Humana, the results document that when older spouses are quasi-randomly induced to join Humana, their younger partners follow. That pattern suggests a causal relationship rather than a positive correlation in unobservable preferences or needs.

Together, the results 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 that seems 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, a key reason for this is that many beneficiaries – especially those with very high costs – have the same individual- and household-cost-minimizing plans. Among beneficiaries who do not have the same individual- and household-cost-minimizing plans, larger savings could be achieved by choosing separate plans.

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 expected out-of-pocket costs. We first demonstrate the same result. We then show that this pattern weakens when we include an indicator variable for one’s spouse being on that plan, though a gap remains. That suggests that, at least on average, spousal following cannot fully explain 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). In addition, the paper is related to the work of Sorensen (2006), who shows that new employees at the University of California tend to choose health insurance plans that are popular among their new colleagues. 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 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 from 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

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

the plans. These insurance companies must offer a plan at or above 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.” 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 [A](#) 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, 72.6 percent of couples in the data are on the same Part D plan. 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. When the older spouse’s drug spending is above the median, only 70 percent of couples pick the same plan, as compared to 75 percent when the older spouse’s spending is below the median. 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.⁵

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 percent of placebo spouses on the same plan represents the rate we would expect if plan choices were solely driven by observable characteristics and the local availability of the plan. 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. The share of recipients on the same plan as their spouse is higher than 66 percent for recipients joining Medicare throughout the years in our sample (Appendix B).

⁵Appendix B presents additional descriptive statistics that document spousal following across income, drug spending, health status, the age difference of the spouses, and year.

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.⁶

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. 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.

Across all specifications that include the “spouse on plan” variable, the coefficient on that variable is large and statistically significant, indicating high value for being on the same plan as one’s spouse. There are several ways to interpret the magnitude of the coefficient. First, we can compare the “spouse-on-plan” coefficient to the coefficient on the plan premium. Doing so provides one estimate of willingness-to-pay to be on the same plan as one’s partner. Such a calculation based on the model in column (6) of Table 2 suggests a large dollar value: $3.4507/0.002143 \approx 1,610$ dollars.

⁶To estimate counterfactual out-of-pocket spending, we estimate the cost of every prescription a recipient fills in every plan available to them 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 (Abaluck and Gruber, 2016). Appendix C describes this process.

Such an exercise, however, is only a valid calculation of willingness-to-pay to the extent that the premium coefficient (and the spouse-on-plan coefficient) is identified. It is possible that the coefficient on premiums is biased, because we do not possess an instrument for premiums in this setting. There may be unobserved plan characteristics that are valued by consumers and also correlated with premiums. As a result, the estimated willingness-to-pay would also be biased.

An alternative method to calculate willingness-to-pay is to scale the “spouse on plan” coefficient by prior estimates of the sensitivity to premiums. The prior literature on Medicare Part D plan choice estimates higher sensitivity to premiums, suggesting that our estimate of willingness-to-pay is biased upward.⁷ Scaling the “spouse on plan” coefficient from the sixth column of Table 2 by those prior estimates yields a willingness-to-pay to be on the same plan as one’s spouse that ranges from \$451 to \$693. That calculation, however, is speculative, because premium coefficients across different discrete-choice models are only comparable up to an unknown scale factor.

Finally, another method to interpret the size of the “spouse on plan” coefficient is to compare it to the magnitude of the coefficient on inertia. The “spouse on plan” coefficient is roughly 77 percent of the magnitude of the inertia coefficient. Polyakova (2016) estimates that beneficiaries are willing to pay \$1,159 to remain on their prior plan. That baseline then suggests a willingness to pay to be on the same plan as one’s spouse of $0.77 \times \$1,159 = \892.43 .

We view these approaches to measuring willingness-to-pay as complementary, suggesting that the true value of willingness-to-pay to be on the same plan as your spouse is likely in the range of \$400 to \$900, rather than in the \$1,500 range suggested by our estimates alone. While this range is large, the numbers are all economically significant. For instance, the estimates of willingness-to-pay above are roughly a third to a half of the average total out-of-pocket costs and premiums in the program.⁸

One additional 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 premium and out-of-pocket-cost coefficients decreases by over 40 percent.

⁷Polyakova (2016) estimates a premium coefficient of -0.00498 (Table 4), for Ho et al. (2017) it is -0.00644 (Appendix Table 11), and for Heiss et al. (2021) it is -0.00765 (Table 10).

⁸Ketcham et al. (2015) estimate average total out-of-pocket costs and premiums of \$1,492 a year in 2010.

That pattern may be consistent with omitted-variables bias. The gap between the coefficient on premiums and the coefficient on out-of-pocket costs involves two mechanisms. First, there is the general over-weighting of premiums focused on by the prior literature. Second, older spouses may overweight premiums in their initial plan choice, and then their younger spouses choose the same plan. In that sense, part of the gap between premiums and out-of-pocket costs seems to be driven by the older spouse’s initial choice.

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 low premiums (\$14.80 each month) and low copayments. As a result, the plan immediately proved popular. Then, in 2014, Humana renamed the 2011 plan the “Humana Preferred RX Plan” and created a new, enhanced Walmart 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.

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 more likely to be female.⁹ 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 [Grembi et al. \(2016\)](#) and [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

⁹Appendix [D](#) provides supporting evidence.

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. 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.

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 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.¹⁰ 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.

4 Overspending Due to Spousal Following

This section presents a decomposition to assess how much spousal following 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, \$540, is larger than estimates in the previous literature, though of the same magnitude.¹¹

¹⁰Appendix D offers additional details and results related to the RD-DD strategy: alternative bandwidths, RD estimates, and placebo analyses.

¹¹The 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.

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 \$464, 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 16 percent in out-of-pocket costs if they chose a different plan than their partner rather than choosing the plan that is best for the couple together. This, however, conceals important heterogeneity. Many spouses have the same optimal plan as each other due to spousal homophily. But 64% of recipients have a different optimal plan than their spouse. The second column of Table 3 performs the same exercise solely for those couples. For such couples, the estimates suggest more to be gained – an additional 24 percent – by choosing separate plans. As an alternative stratification, the last two columns of Table 3 perform the same exercise separately for younger and older spouses. The overspending numbers are similar across younger and older spouses: personal overspending is roughly \$540 in either case.

Table 3 therefore illustrates two implications of spousal following. First, couples’ tendency to choose the same plan explains a small share of total overspending. Almost all of the possible savings in out-of-pocket costs can be achieved by couples choosing a cost-minimizing plan jointly. The difference is somewhat starker, however, for the 64-percent of the sample that has a different optimum plan as their spouse. For that group, spousal following leads to a larger increase in out-of-pocket costs they could otherwise avoid. In that fashion, these results suggest real savings may be possible in settings in which peers are more likely to have different optimal plans.

Second, Table 3 rules out a tempting hypothesis. One might imagine that older spouses choose plans optimally, leading to minimal overspending, and then younger spouses simply follow them, leading to large overspending for them. That is not the case – both older and younger spouses have roughly similar levels of overspending. This could be driven by misoptimization by the older spouse that is then followed by a younger spouse with similar drug needs. It could also be driven by joint misoptimization, in which the couple chooses together at the time that the older spouse enters, but chooses poorly. Ultimately, we cannot differentiate between such potential mechanisms.

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 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 the 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.

It is also possible that there exist efficiencies in terms of the burden of choosing a plan. Consumers may pay significant search costs when choosing a plan ([Brown and Jeon, 2024](#)). And thus perhaps some couples choose the same plan simply because it is easier to learn about and choose one plan once – when the older spouse joins Medicare – than to learn about the options again when the younger spouse enters and choose two plans.

In that sense, there exist a variety of mechanisms that may drive spousal following. Some couples may choose the same plan due to information frictions and search costs and others due to administrative costs. It is also possible that spouses specialize in financial decisions, and only one spouse chooses plans, regardless of who is gaining Medicare coverage. In addition, some forward-looking couples may consider the younger spouse’s future costs when first choosing a plan for the older spouse.

The data do not allow us to distinguish between these mechanisms, making this a potentially fruitful area for future research. That said, the relatively large willingness-to-pay estimates suggest that spousal following is unlikely to reflect a conventional cost–benefit calculation. The most 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. The key issue is 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.
- 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.
- Brown, Zach Y and Jihye Jeon**, “Endogenous information and simplifying insurance choice,” *Econometrica*, 2024, *92* (3), 881–911.
- 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.
- Grembi, Veronica, Tommaso Nannicini, and Ugo Troiano**, “Do Fiscal Rules Matter?,” *American Economic Journal: Applied Economics*, 2016, *8* (3), 1–30.

- 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.
- Ketcham, Jonathan D, Claudio Lucarelli, and Christopher A Powers**, “Paying Attention or Paying Too Much in Medicare Part D,” *American Economic Review*, 2015, 105 (1), 204–233.
- 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.
- 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.
- Sorensen, Alan T.**, “Social Learning and Health Plan Choice,” *Rand Journal of Economics*, 2006, *37* (4), 929–945.
- 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			4.2577*** (0.0136)	4.2456*** (0.0137)	3.6981*** (0.0151)	3.4507*** (0.0157)
Premium (in hundreds)	-0.1706*** (0.0022)	-0.1796*** (0.0023)	-0.1235*** (0.0028)	-0.1331*** (0.0029)	-0.1481*** (0.0034)	-0.2143*** (0.0045)
OOP Cost (in hundreds)		-0.0696*** (0.0024)		-0.0716*** (0.0046)	-0.0663*** (0.0049)	-0.0732*** (0.0073)
On Plan Last Year						4.4537*** (0.0212)
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 4,435,391 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



These figures describe enrollment in Humana plans before and after the introduction of the Humana-Walmart plans. All panels present 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 between 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 control periods. 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 control periods.

Table 3. Overspending Decomposition

	(1)	(2)	(3)	(4)
	Full	Different	Younger	Older
	Sample	Plan	Spouses	Spouses
Overspending relative to <i>personal</i> cost-minimizing plan	539.96	610.32	534.84	545.08
Overspending relative to <i>household</i> cost-minimizing plan	464.41	493.10	457.19	471.62
Percent of sample	100.00	64.00	50.00	50.00

This table presents average overspending for Part D recipients in California, Florida, New York, Pennsylvania, and Texas. The first row presents the average 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 difference between each individual's out-of-pocket spending in their chosen plan and the plan in which, if chosen by them and their spouse, they would have incurred the smallest out-of-pocket costs as a household. The first column presents statistics for the entire sample, and the second column presents statistics for those whose cost-minimizing plan does not match their partner's cost-minimizing plan. The remaining columns present statistics separately for younger spouses and older spouses.

A 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 – spousal following – 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 [A1](#) 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 A1. 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.

B Appendix: Additional Descriptive Results

This appendix section presents additional descriptive results on couples' plans choice, including by ZIP code income, drug spending level, health status, and age difference.

As an illustration of how couples sort across plans, Appendix Figure B1 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.


Appendix Figure B2 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.

Appendix Figures B3 and B4 plot the same outcomes by deciles of drug spending and number of same chronic conditions.

Appendix Figure B5 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.


Appendix Figure B6 plots the share of recipients on the same drug plan as their spouse by year. 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 throughout those years.

Appendix Figure B1. 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 B2. 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 B3. 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 B4. 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 B5. 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 B6. Spousal Following Separately by Year



This figure plots the share of Medicare recipients who choose the same plan as their spouse by year.

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 \(2011\)](#) and [Abaluck and Gruber \(2016\)](#). 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 spouses who enter Medicare in different years than each other. Finally, we limit the sample to 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.

D Appendix: Additional Regression-Discontinuity Results

This appendix section presents additional figures and tables supporting the regression-discontinuity difference-in-difference analysis.


Appendix Figures D1 and D2 illustrate that because 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 more likely to be female.

Appendix Tables D1 and D2 present estimates of the main regression-discontinuity difference-in-difference specification with different bandwidths. The main estimates change little across varying bandwidths.

Appendix Figure D3 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 D4 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.


Appendix Figure D5 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.

Appendix Figure D1. Humana Enrollment 2011 Case Study: Controls




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 D2. 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 D3. 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 D4. 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 D5. 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.

Appendix Table D1. Bandwidth Selection: Humana Enrollment 2011 Case Study

Bandwidth	30 Days	60 Days	90 Days	120 Days	150 Days
Younger spouse on Humana	0.0531*** (0.0104)	0.0544*** (0.0091)	0.0524*** (0.0085)	0.0512*** (0.0078)	0.0497*** (0.0070)
<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 D2. Bandwidth Selection: Humana Enrollment 2014 Case Study

Bandwidth	30 Days	60 Days	90 Days	120 Days	150 Days
Younger spouse on Humana	0.0612*** (0.0094)	0.0617*** (0.0062)	0.0590*** (0.0059)	0.0571*** (0.0056)	0.0553*** (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 $p < 0.01$.