Appendix: An illustrative example of cream-skimming incentives

In this appendix we use a simple example to illustrate how cream-skimming incentives can still exist in the presence of a "perfect" risk score under a given contract when individuals are heterogeneous in their behavioral responses to contracts. We also consider the optimal risk adjustment policy for the government in the presence of such incentives.

For concreteness, we consider a specific, highly-stylized environment. Our objective is not to derive results that can be directly applied to a specific real-world context, but simply to help illustrate the potential new considerations that come into play as a result of the richer heterogeneity documented in the paper.

Setting We assume the government offers a default contract, and consider a private (monopolist) insurer who offers a contract that competes to attract beneficiaries from the default contract. The government reimburses the private insurer based on the risk scores of the beneficiaries it attracts. The key point is that an individual's medical spending may differ under the private contract and the government default contract. We show how this impacts provider cream-skimming incentives under a given set of reimbursement rules.

For simplicity, we assume that the default public coverage provides full insurance (i.e. $c = 0$ in the framework of Section III). As a result, in our framework – see especially equations (2) and (3) – beneficiary i , who can be fully described by his two-dimensional type (λ_i, ω_i), chooses medical spending level $\lambda_i + \omega_i$ and obtains utility $u_i^* = y_i + \omega_i/2$. The associated government spending g_i is then $\lambda_i + \omega_i$. These assumptions are summarized in first three rows of Appendix Table A2, in the first column, to make them more easily comparable to our assumptions about the private plan, summarized in the second column.

We make the (extreme) assumption that the private plan can completely eliminate ω-related medical spending. In other words, they can incentivize physicians to "perfectly" distinguish medical spending associated with λ_i – that is, medical spending that would have been made by the beneficiary even with no coverage – from medical spending associated with ω_i, which would not have occurred if the beneficiary was uninsured. Moreover, we assume that the private plan will only cover λ_i-related medical spending. Under the private plan therefore, medical spending associated with beneficiary i is λ_i and he obtains utility $u_i^* = y_i$ as summarized in Appendix Table A2.

The government reimburses the private insurer based on the current Medicare risk scores of the beneficiaries it enrolls; we denote these risk scores by r_i. We assume that the government can only observe medical spending under its own, public contract, and thus can only score beneficiaries based on their predicted medical spending under the public contract. Risk scoring is thus based on a prediction model of medical spending under public coverage. Specifically, the risk score r_i is medical spending under the public plan; $r_i = \lambda_i + \omega_i$. As Figures 3 and 4 illustrated empirically, this risk score does not distinguish between beneficiary costs arising from λ or from ω. In keeping with our focus on a challenge to risk scoring that exists even if risk scores are "perfect" in the statistical sense, we also assume here that there are no residual characteristics of the individual that predict $\lambda_i + \omega_i$ conditional on r_i.

Given this "perfect" risk score r_i, private insurers receive a risk-adjusted transfer from the government, $g(r_i)$, for covering beneficiary i. It will be convenient to define $g(r_i) \equiv r_i + s(r_i)$ as the sum of the cost this beneficiary would have incurred under the public plan, r_i, and an additional (positive or negative) subsidy $s(r_i)$. Therefore, as shown in row 3 of Appendix Table A2, government spending g_i under the public plan is given by beneficiary medical spending under this plan $\lambda_i + \omega_i$, and under the private plan it is given by $g(r_i) = \lambda_i + \omega_i + s(\lambda_i + \omega_i)$.

As shown in row 4, insurer profits π_i from covering individual i are given by the government (risk-adjusted) transfer $g(r_i)$ minus the cost to the insurer of covering individual i under the offered contract; private provider costs are λ_i by assumption.[7]

The bottom row of Appendix Table A2 shows the implications of each insurance allocation for total surplus associated with beneficiary i. We define total surplus as the sum of consumer surplus and producer surplus, minus government spending (and its associated costs). For individual i it is given by

$$TS_i = u_i^* + \pi_i - (1 + k)g_i, \tag{A1}$$

[7] We abstract from any administrative costs of the private insurer which will not affect the fundamental selection analysis.

where k is the shadow cost of public funds.

Given our assumptions, it is socially efficient for everyone to be covered by the private plan if there is no subsidy ($s(r_i) = 0$), with the gain in efficiency from enrolling beneficiary i in the private plan increasing in ω_i. This is because we have assumed that under the private plan, insurance coverage is still full, but does not apply to ω_i-related medical spending; and under the assumptions of our model, ω_i-related medical spending is socially inefficient. As a result, total surplus under the private plan is higher by $0.5\omega_i$ for any beneficiary i (see row 5 of Appendix Table A2). However, as we will see below, the private provider would not find it optimal to cover all beneficiaries due to cream-skimming incentives, so an additional government subsidy $s(r_i)$ would be needed to get more beneficiaries enrolled in the private plan. Optimal risk adjustment would thus trade off these two offsetting forces: increasing allocative efficiency from inducing more individuals to be covered by private plan against the social cost of public funds for the subsidy needed to enroll additional individuals in the private plan.

Private provider optimization problem We now consider the impact of the foregoing set-up for the provider's optimal (profit-maximizing) contracts and, therefore, for equilibrium cream-skimming. We model the monopolist private insurer's problem as a standard optimal contracting model with incomplete information.

The private provider observes the risk score r (which, recall, is simply spending under the public contract) and offers a family of contracts that are a function of the observed risk score r and a self-reported (by beneficiaries) type ω'. If a beneficiary chooses a contract $p(r, \omega')$, he would pay a premium $p(r, \omega')$ to the private insurer, and the private insurer would cover medical spending of $\lambda' = r - \omega'$; we note that, under our assumptions, this is the efficient amount of medical spending for someone with risk score r and type ω'. We assume that people know their true type (λ, ω) when choosing insurance plans. Insurers then design the contracts such that each beneficiary truthfully reveals his type, thus allowing the private provider to only authorize the λ-related medical spending associated with the true λ.

Consider the utility of beneficiary of type (λ, ω) from a private contract $p(r, \omega')$. Recall that $r = \lambda + \omega$ is observed, and that individual medical spending under the private contract would be given by $\lambda' = r - \omega' = \lambda + \omega - \omega'$. Plugging this into the utility function in equation (1) we obtain utility under the private contract:

$$u(\lambda, \omega; \omega') = \left[(\omega - \omega') - \frac{1}{2\omega} (\omega - \omega')^2 \right] + y - p(r, \omega'). \tag{A2}$$

By constrast, staying in the public plan would result in utility of $y + 0.5\omega$, as derived earlier (see Appendix Table A2).

Because r is observed and contractible, we can then solve the profit maximization problem of the private provider separately for each r. We denote by $F_{\omega|r}$ the conditional (on r) cumulative distribution function of ω. The private provider's problem is to choose the menu of contracts $p(r, \omega')$ in order to maximize profits:

$$\max_{p(r,\omega)} \pi(r) = \int [p(r, x) + g(r) - (r - x)] \, dF_{\omega|r}(x) = \int [p(r, x) + s(r) + x] \, dF_{\omega|r}(x), \tag{A3}$$

subject to an incentive compatibility (IC) constraint that makes beneficiaries choose the contract that matches their type

$$u(\lambda, \omega; \omega) = y - p(r, \omega) \geq u(\lambda, \omega; \omega') \; \forall \omega', \tag{A4}$$

and an individual rationality (IR) constraint that makes beneficiaries willing to opt out of the public plan and instead enroll in private coverage

$$u(\lambda, \omega; \omega) = y - p(r, \omega) \geq y + \omega/2. \tag{A5}$$

The IC constraint can be written as $\omega = \arg\max_{\omega'} u(\lambda, \omega; \omega')$. A necessary and sufficient condition is that the IC constraint's first order condition holds. Solving the IC constraint using utility from the private contract defined in equation (A2), gives $-1 - \partial p / \partial \omega = 0$, implying that $p(r, \omega) = t(r) - \omega$, where $t(r)$ is the integration constant (which could depend on r, as the solution is conditional on r). Substituting this schedule into the IR constraint above (equation (A5)), we obtain $y - (t(r) - \omega) \geq y + \omega/2$. Thus selection into private coverage is given by:

$$\omega \geq 2t(r) \tag{A6}$$

Equation (A6) describes equilibrium selection under the profit maximizing contract: for every risk score r, higher ω beneficiaries select into the private contract while lower ω beneficiaries remain in the public plan. Thus, on one hand, selection will be in general favorable: the beneficiaries for whom it is most socially efficient to be covered by the private provider will be covered by the private provider. On the other hand, because risk scoring does not capture this second dimension of heterogeneity, some fraction of beneficiaries will inefficiently remain covered by the public plan.

Given the equilibrium selection rule in equation (A6), the profit maximization problem from equation (A3) becomes

$$\max_{t(r)} \pi(r) = (t(r) + s(r)) \Pr(\omega \geq 2t(r)|r) = (t(r) + s(r)) \left[1 - F_{\omega|r}(2t(r))\right]. \tag{A7}$$

The monopolist therefore sets $t^*(r)$ to solve the first order condition

$$t^*(r) = \frac{1 - F_{\omega|r}(2t^*(r))}{2f_{\omega|r}(2t^*(r))} - s(r), \tag{A8}$$

with $f_{\omega|r}(x) = F'_{\omega|r}(x)$.

This is a familiar profit maximization problem, very similar to the textbook optimal pricing problem for a monopolist facing a downward sloping unit demand curve. The profit function in equation (A7) has the familiar form of $\pi = (p - c)D(p)$ where, here, price p is given by $t(r)$, marginal cost c is given by $-s(r)$, and demand $D(p)$ is given by $1 - F$. As in the textbook case, the monopolist trades off price vs. quantity: raising the price $t(r)$ will result in greater profits on inframarginal beneficiaries, but a loss of marginal beneficiaries to the public plan. The extent of the loss of marginal beneficiaries (and hence the optimal price $t(r)$) depends on the shape of the demand curve, which in this case depends on the hazard rate of the distribution of ω. The problem becomes similar to the textbook monopolist pricing problem because the private provider does not observe ω, and his cream-skimming incentive – the greater profit he obtains from higher-ω beneficiaries – is exactly offset by the increased incentive of the higher-ω beneficiaries to remain in the public plan.

Implications for designing risk adjustment The analogy to the monopolist pricing problem also makes it easy to see the role that alternative risk adjustment formulations could play. Looking at equation (A7), the government subsidy $s(r)$ can be thought of as shifting the monopolist's marginal cost, since $s(r)$ enters the profit function just like the negative of marginal cost. This yields clear and natural comparative statics: an increase (decrease) in the subsidy $s(r)$ would provide more (less) powerful incentives to the private provider to enroll additional beneficiaries by reducing (increasing) the "unit price" $t(r)$. Therefore, as can be seen directly in equation (A8), changes in the subsidy $s(r)$ are partially passed through to the premium $p^*(r,\omega) = t^*(r) - \omega$. The government subsidy affects the private provider's profit maximizing pricing, and thereby affects the equilibrium selection of individuals to the private plan.

Absent any social cost of public funds, the bottom row of Appendix Table A2 makes it clear that the optimal subsidy should be high enough, so that the private provider would set $t(r) = 0$ and thereby, by equation (A6), enroll all beneficiaries. However, when the social cost of public funds is positive ($k > 0$), the optimal subsidy is set to resolve a simple tradeoff: higher $s(r)$ would efficiently enroll more beneficiaries under the private plan, but would be associated with greater costs of public funds.

To see this tradeoff more formally, consider the government's optimization problem. For a given risk score r, the government chooses the optimal transfer $g^*(r) = r + s^*(r)$ in order to maximize total surplus subject to the private insurer setting $t^*(s(r);r)$ optimally.

Relative to the total surplus that would arise from covering all beneficiaries in the public plan, the incremental surplus from allocating an individual to the private plan instead is then given by

$$\Delta TS(s(r);r) = \left[\frac{1}{2}E\left[\omega|\omega \geq F_{\omega|r}(2t^*(s(r);r))\right] - ks(r)\right]\left[1 - F_{\omega|r}(2t^*(s(r);r))\right]. \qquad \text{(A9)}$$

By increasing $s(r)$ the government indirectly decreases $t^*(r)$ and thus (socially efficiently) enrolls more beneficiaries in the private plan, at the cost of increasing the cost of public funds for all inframarginal beneficiaries who were already enrolled in the private plan. For a given risk score r, the government would optimally set $s(r)$ to maximize $\Delta TS(s(r);r)$, with the optimal subsidy $s(r)$ decreasing in the cost of public funds k.

Thus far we have considered the optimal government subsidy $s^*(r)$ for a given r. Our analysis has shown that the optimal risk adjustment – the function $g(r)$, or equivalently the subsidy $s(r) = g(r) - r$ – can be solved for each r separately. Analysis of optimal risk adjustment requires determining how the optimal subsidy $s^*(r)$ varies with r. This is in the spirit of inquiry of Glazer and McGuire (2000), who found (in their unidimensional heterogeneity model) that the optimal risk adjustment should amplify the observed risk scores.

A full characterization of $s(r)$ would require assumptions (or ideally evidence) about the specific objects, most importantly about the distribution $F_{\omega|r}$, which determines the shape of the demand curve, and the social costs of public funds k, which determines the cost of subsidies designed to increase the set of people (efficiently) covered by the private provider.[8] The social cost of public funds is typically assumed to be about 0.3 (Poterba, 1996). The conditional distribution $F_{\omega|r}$ – i.e. the shape of unobserved (by the government) individual type conditional on observed individual spending under the government contract – would need to be estimated in the specific application. In our empirical context, for example, some evidence on the shape of this function was provided by our analysis of who "bunches" at the kink in the donut hole. Figure 3 indicated that healthier individuals were more likely to bunch at the kink, suggesting that ω and λ are negatively correlated, at least around the donut hole, which may provide some guidance regarding $F_{\omega|r}$ (recall, in our example $r = \omega + \lambda$).

[8] Under a given set of assumptions about k and $F_{\omega|r}(\cdot)$, it is easy (although typically cumbersome) to solve for the optimal $s(r)$. For example, inspection of equation (A9) indicates that if $\omega|r$ is uniformly distributed over $[0, A(r)]$, the optimal subsidy $s^*(r)$ would scale proportionally with $A(r)$. For the general case, however, the optimal $s(r)$ could be either positive or negative, and could either increase or decrease (or not even be monotone) in r.

Figure 1: Medicare Part D standard benefit design (in 2008)

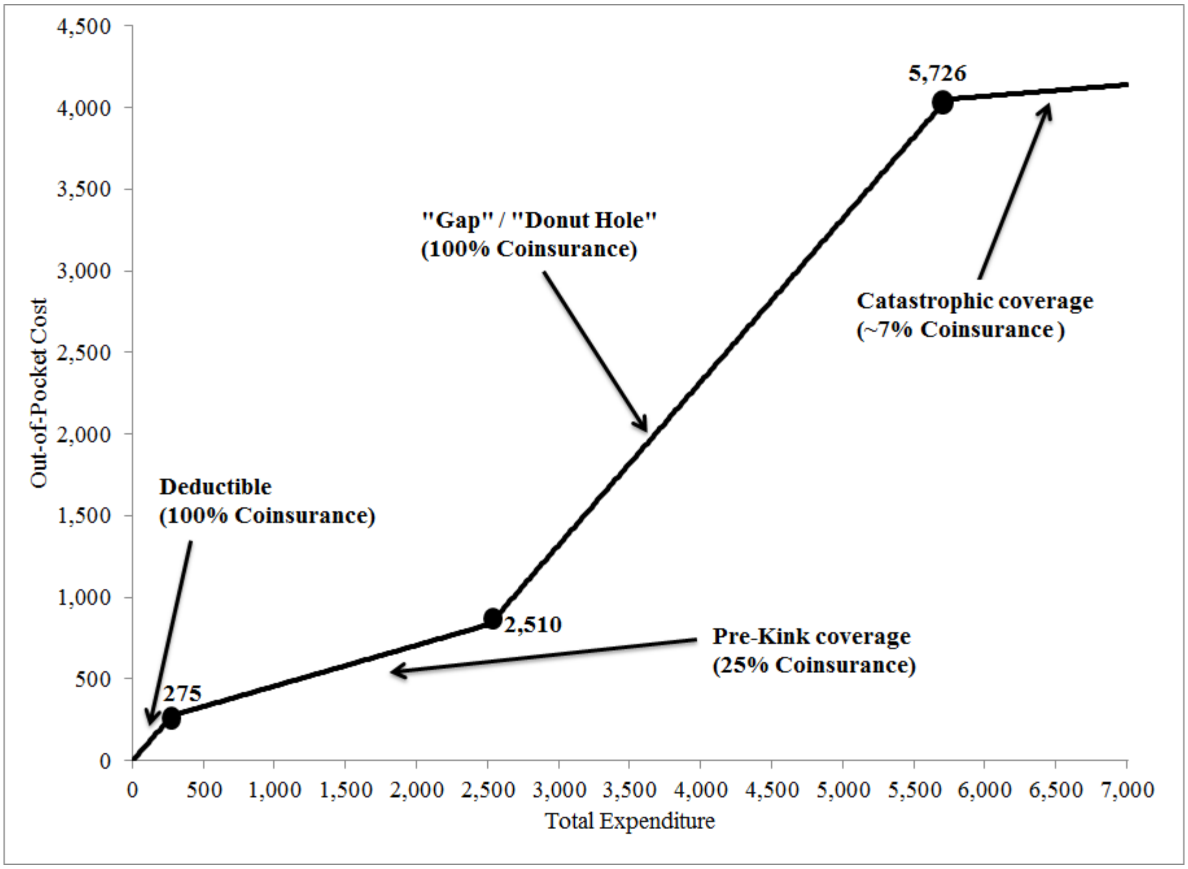

The figure shows the standard benefit design in 2008. "Pre-Kink coverage" refers to coverage prior to the Initial Coverage Limit (ICL) which is where there is a kink in the budget set and the gap, or donut hole, begins. The level at which catastrophic coverage kicks in is defined in terms of out-of-pocket spending (of $4,050), which we convert to the total expenditure amount provided in the figure. Once catastrophic coverage kicks in, the actual standard coverage specifies a set of co-pays (dollar amounts) for particular types of drugs; in the figure we use show a 7 percent co-insurance rate, which is the empirical average of these co-pays in our data.

Figure 2: Bunching of annual spending around the kink

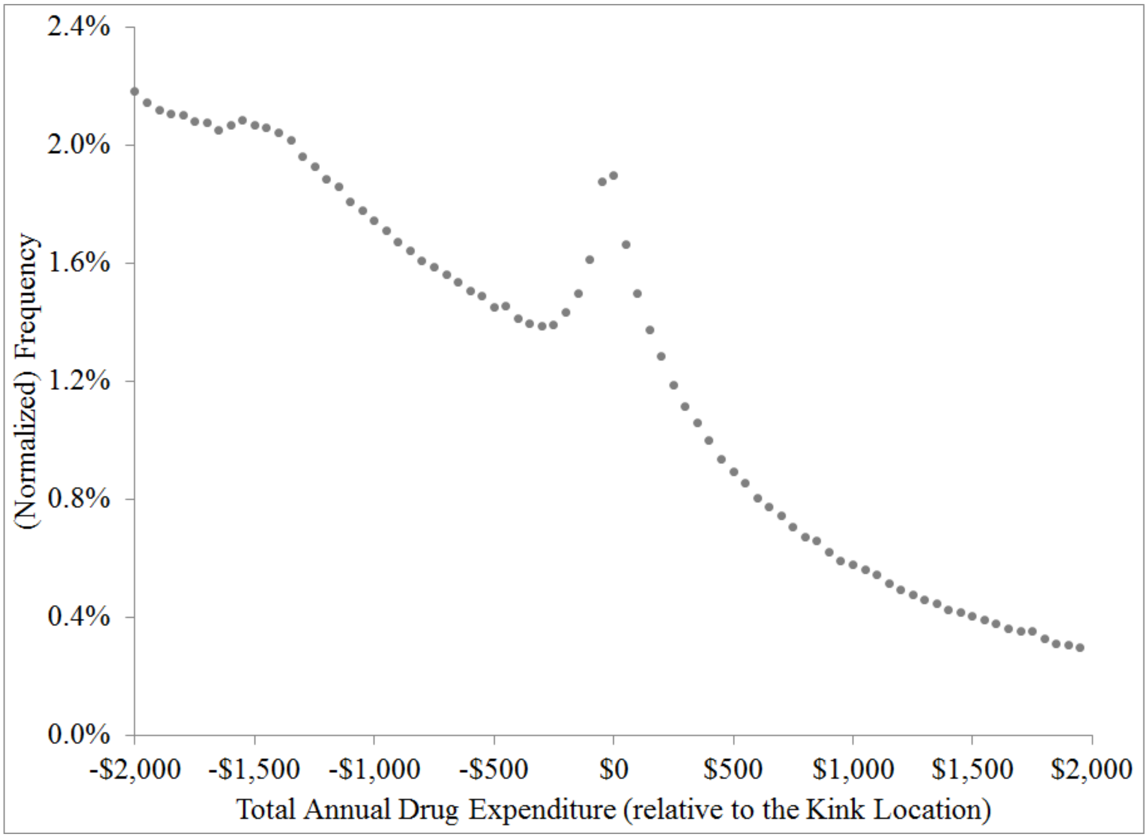

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within $2,000 of the (year-specific) kink location. The points in the figure display the distribution of annual spending; each point represents the set of people that spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. We normalize the frequencies so that they add up to one for the range of annual spending shown. N = 2,506,305.

Figure 3(a): Variation in demographics around the kink

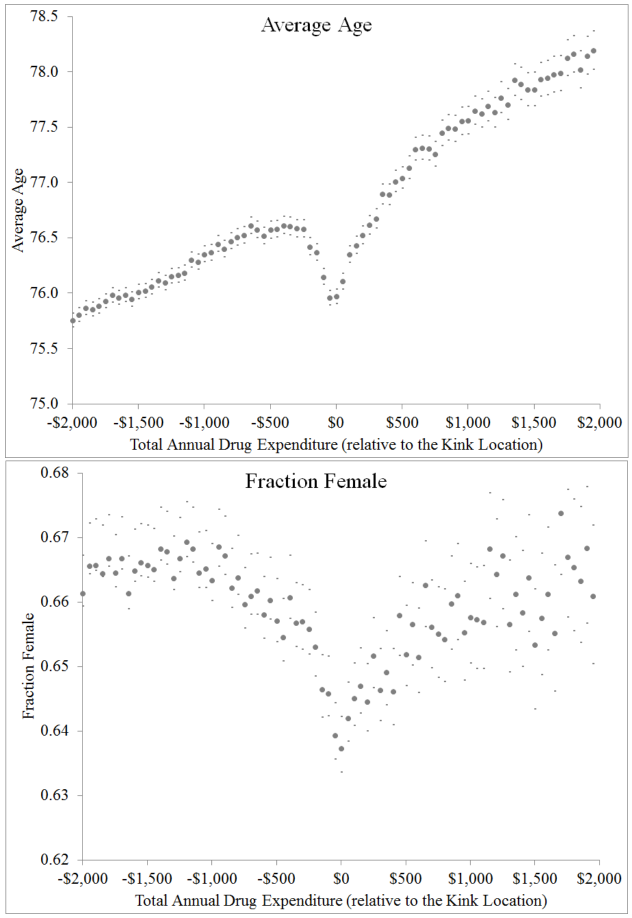

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within $2,000 of the (year-specific) kink location. The points in the figure display the statistic described on the y-axis for each group of beneficiaries who spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. The small lines above and below the points represent the 95 percent confidence interval for each point. N = 2,506,305.

Figure 3(b): Variation in selected health conditions around the kink

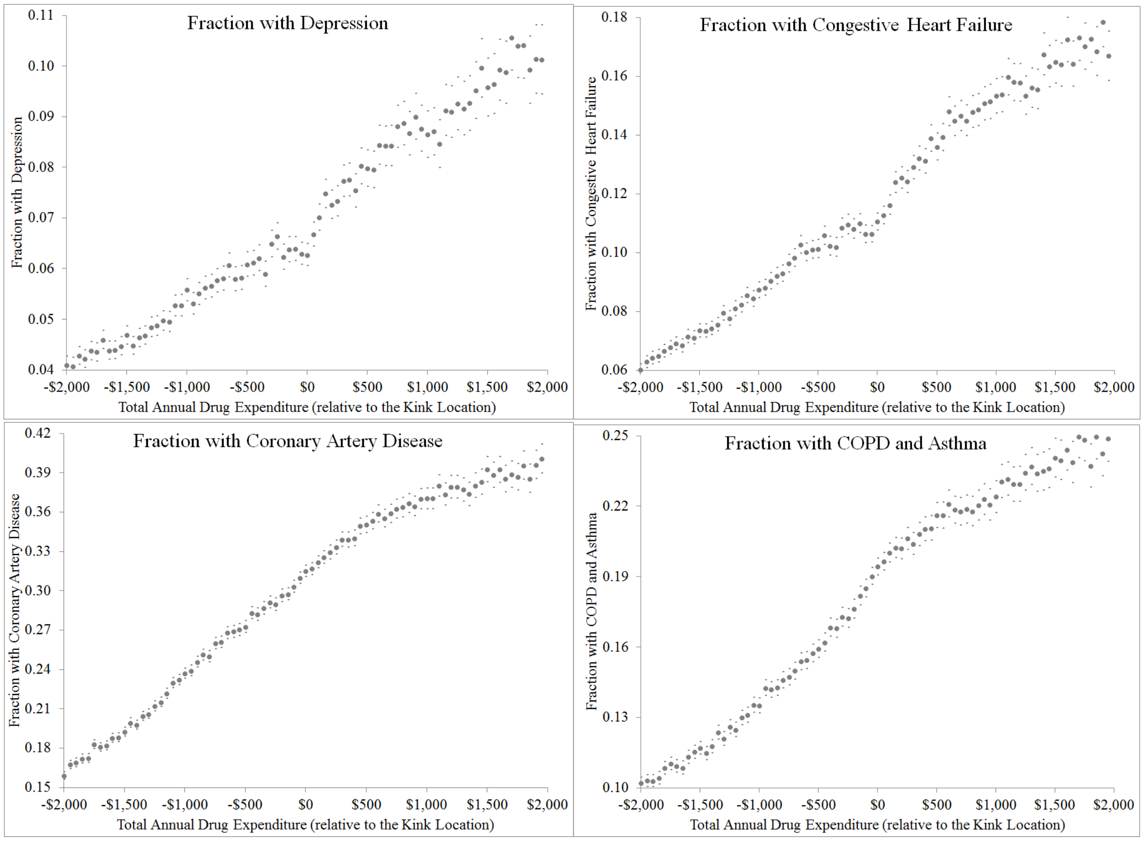

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within $2,000 of the (year-specific) kink location. The points in the figure display the statistic (for the subsequent coverage year, so covering the years 2008-2010) described on the y-axis for each group of beneficiaries who spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. The small dots above and below the points represent the 95 percent confidence interval for each point. N = 2,506,305.

Figure 3(c): Variation in subsequent mortality and healthcare utilization around the kink

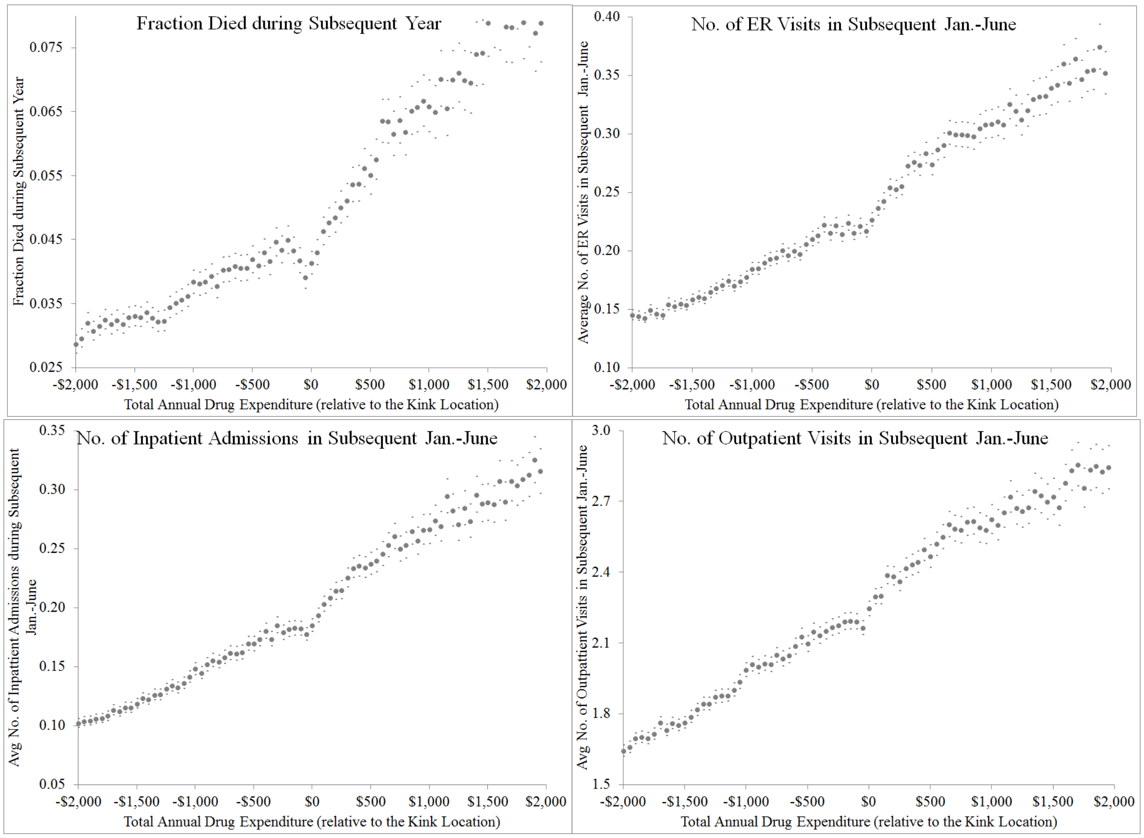

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within $2,000 of the (year-specific) kink location. The points in the figure display the statistic (for the subsequent coverage year, so covering the years 2008-2010) described on the y-axis for each group of beneficiaries who spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. The small dots above and below the points represent the 95 percent confidence interval for each point. N = 2,506,305.

Figure 4(a): Variation in "endogenous" risk score around the kink

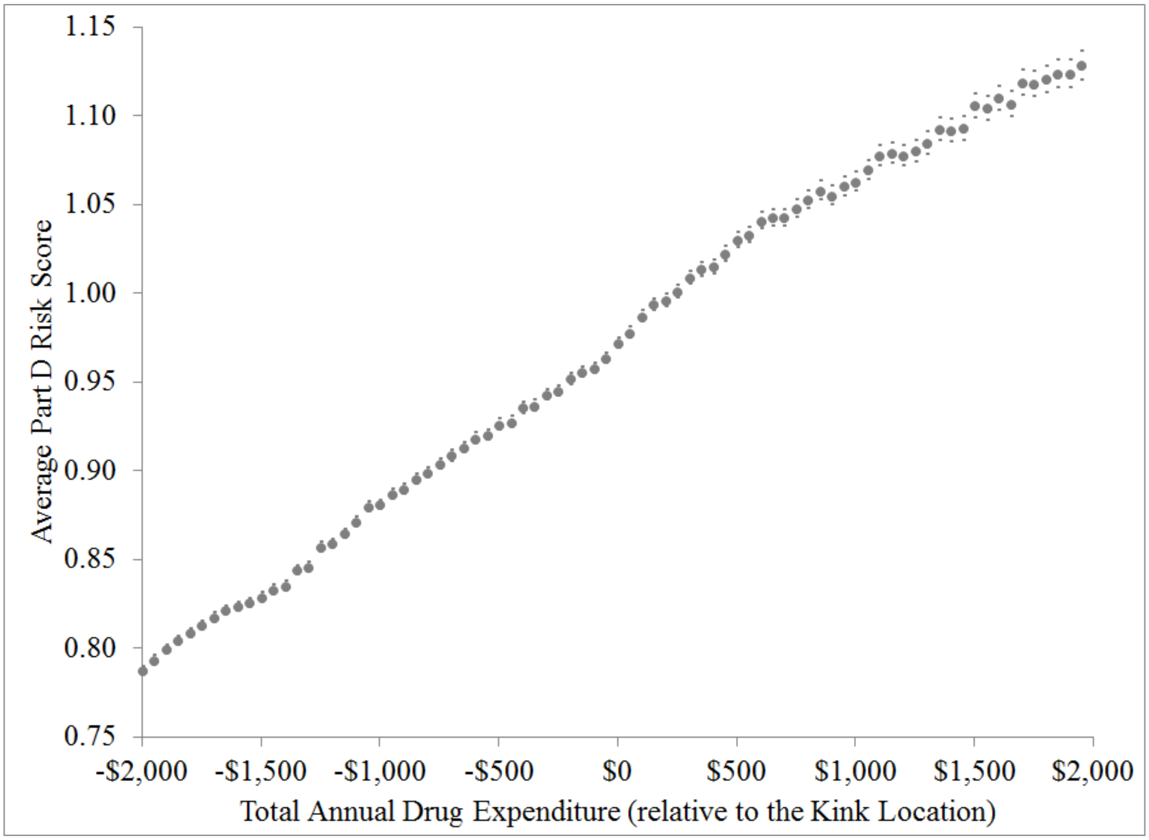

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within $2,000 of the (year-specific) kink location. The points in the figure display the statistic described on the y-axis for each group of beneficiaries who spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. The small dots above and below the points represent the 95 percent confidence interval for each point. N = 2,506,305.

Figure 4(b): Magnitude of offsetting effects

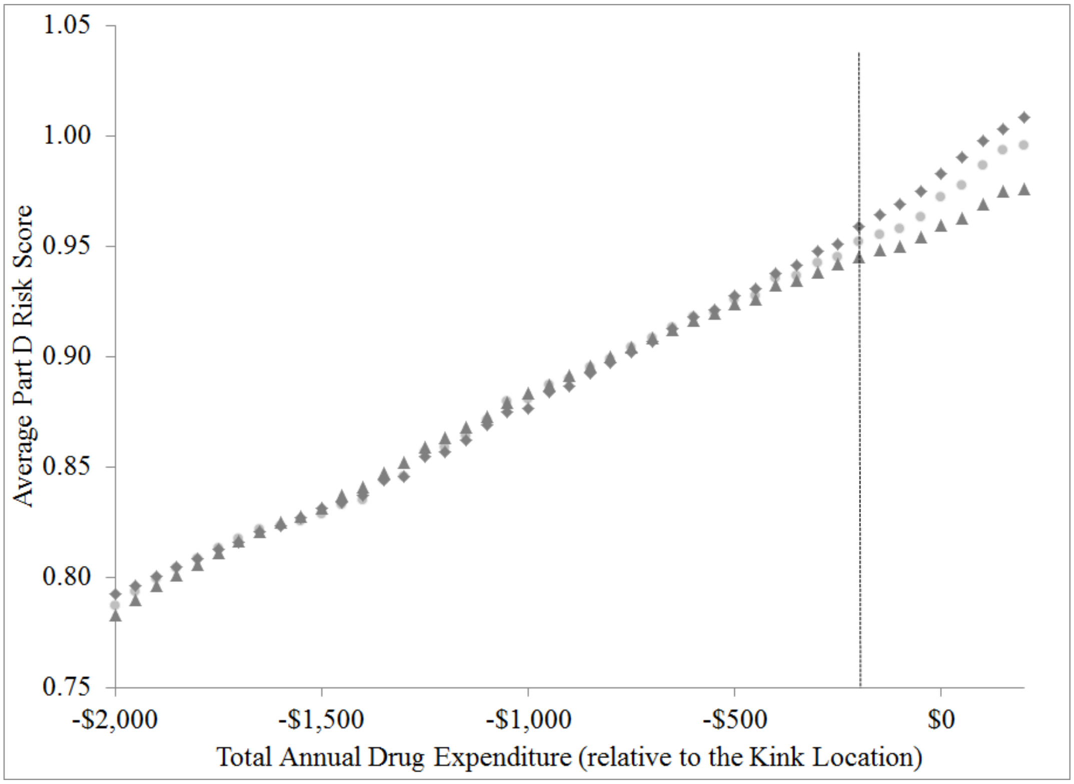

Total annual prescription drug spending on the x-axis is reported relative to the (year-specific) location of the kink, which is normalized to zero. Sample uses beneficiary-years in our 2007-2009 baseline sample whose annual spending is within -$2,000 to +$200 of the (year-specific) kink location. The points in the figure display the statistic described on the y-axis for each group of beneficiaries who spent up to $50 above the value that is on the x-axis, so that the first point represents individuals who spent between -$2,000 and -$1,950 from the kink, the second point represents individuals between -$1,950 and -$1,900, and so on. N = 1,948,900. The light gray series presents the actual risk score, replicating Figure 4(a). The top (grey squares) series shows the risk score generated by taking the predicted values of the risk score components that exhibit dips at the kink, and the actual values for the rest. The bottom (grey triangles) series shows the risk score generated by taking the predicted values of the risk score components that exhibit bunching at the kink, and the actual values for the rest. The predictions for each component of the risk score is generated by fitting a trend in that component for spending more than $200 under the kink (see text for more details).

Table 1: Components underlying the effect on risk scores around the kink

	Incidence around the kink			Share of Risk-
	Actual	"Predicted"	Difference	Score difference
Top 10 components with positive kink incidence				
Chronic Obstructive Pulmonary Disease and Asthma	0.1908	0.1784	0.0124	20.50%
Diabetes with Complications	0.0908	0.0816	0.0091	19.13%
Breast, Lung, and Other Cancers and Tumors	0.0582	0.0520	0.0062	10.72%
Alzheimer's Disease	0.0203	0.0179	0.0024	9.32%
Diabetes without Complications	0.2020	0.1962	0.0058	8.45%
Esophageal Reflux and Other Disorders of Esophagus	0.2146	0.2082	0.0064	7.26%
Inflammatory Bowel Disease	0.0107	0.0093	0.0014	3.21%
Diabetic Retinopathy	0.0278	0.0237	0.0041	3.20%
Parkinson's Disease	0.0127	0.0119	0.0009	3.06%
Major Depression	0.0196	0.0187	0.0010	2.26%
Top 10 components with negative kink incidence				
Hypertension	0.6531	0.6735	0.0203	33.48%
Disorders of Lipoid Metabolism	0.7344	0.7530	0.0186	21.63%
Osteoporosis, Vertebral and Pathological Fractures	0.1730	0.1874	0.0144	13.13%
Open-Angle Glaucoma	0.0918	0.0999	0.0081	11.28%
Atrial Arrhythmias	0.1361	0.1460	0.0099	5.99%
Congestive Heart Failure	0.1117	0.1148	0.0031	5.40%
Thyroid Disorders	0.2525	0.2596	0.0071	2.64%
Coronary Artery Disease	0.3107	0.3116	0.0009	1.23%
Depression	0.0659	0.0668	0.0009	1.20%
Cereborvascular Disease, Except Hemorrhage or Aneurysm	0.1513	0.1522	0.0009	0.59%

Table presents the top 10 components responsible for the positive and negative risk score effect presented in Figure 4(b). The first column ("Actual") reports the average value of the component around the kink (specifically, between -$200 and $200 of the kink, whose value is normalized to 0). The second column ("Predicted") reports the average predicted value of each component, by extrapolating a linear relationship from the (-$2,000,-$200) range. The third column ("Difference") reports the difference between Actual and Predicted. The last column ("Share of Risk-Score difference") reports the share of each component in generating the positive (top panel) and negative (bottom panel) risk score effect presented in Figure 4(b). This is computed by multiplying the difference associated with each component (as reported in the third column) by the risk-score coefficient on that component in the risk-adjustment formula, and normalizing this product by the sum of all these products that are associated with positive (top panel) and negative (bottom panel) deviations around the gap.

Appendix Table A1: Risk type ranking under alternative contract designs

Spending under standard contract		Predicted spending under counterfactual contract				
Percentile range (1)	Average spending (2)	Move >1 bin down (3)	Move 1 bin down (4)	Remain in same bin (5)	Move 1 bin up (6)	Move >1 bin up (7)
Panel A. Counterfactual contract "fills" the coverage gap						
0-12[a]	0	--	--	0.99	0.00	0.00
12-15[a]	15	--	0.00	0.99	0.00	0.00
15-20	65	0.00	0.01	0.98	0.00	0.00
20-25	227	0.00	0.01	0.98	0.01	0.00
25-30	527	0.00	0.02	0.96	0.01	0.01
30-35	760	0.00	0.03	0.92	0.02	0.02
35-40	950	0.00	0.07	0.87	0.04	0.03
40-45	1,128	0.00	0.10	0.83	0.04	0.04
45-50	1,305	0.00	0.13	0.78	0.04	0.04
50-55	1,489	0.00	0.17	0.71	0.06	0.06
55-60	1,691	0.00	0.22	0.65	0.06	0.07
60-65	1,916	0.00	0.27	0.57	0.07	0.09
65-70	2,162	0.00	0.34	0.45	0.10	0.12
70-75	2,404	0.00	0.43	0.32	0.13	0.12
75-80	2,590	0.00	0.43	0.32	0.17	0.09
80-85	2,843	0.11	0.39	0.29	0.14	0.07
85-90	3,296	0.07	0.40	0.36	0.13	0.04
90-95	4,140	0.01	0.37	0.52	0.10	--
95-100	7,394	0.00	0.21	0.79	--	--

Spending under standard contract		Predicted spending under counterfactual contract				
Percentile range (1)	Average spending (2)	Move >1 bin down (3)	Move 1 bin down (4)	Remain in same bin (5)	Move 1 bin up (6)	Move >1 bin up (7)
Panel B. Counterfactual contract removes deductible and increases pre-gap cost sharing						
0-12[a]	0	--	--	0.68	0.08	0.24
12-15[a]	15	--	0.50	0.12	0.16	0.21
15-20	65	0.38	0.15	0.22	0.14	0.11
20-25	227	0.27	0.36	0.22	0.09	0.06
25-30	527	0.05	0.38	0.41	0.11	0.05
30-35	760	0.03	0.24	0.58	0.12	0.04
35-40	950	0.03	0.14	0.62	0.18	0.02
40-45	1,128	0.04	0.11	0.61	0.21	0.02
45-50	1,305	0.05	0.11	0.58	0.25	0.02
50-55	1,489	0.06	0.10	0.57	0.26	0.01
55-60	1,691	0.06	0.09	0.58	0.26	0.01
60-65	1,916	0.07	0.08	0.62	0.23	0.00
65-70	2,162	0.06	0.09	0.64	0.21	0.00
70-75	2,404	0.05	0.09	0.66	0.20	0.01
75-80	2,590	0.04	0.10	0.71	0.16	0.00
80-85	2,843	0.05	0.07	0.78	0.09	0.01
85-90	3,296	0.03	0.05	0.86	0.06	0.00
90-95	4,140	0.01	0.05	0.92	0.02	--
95-100	7,394	0.01	0.01	0.98	--	--

Table is based on the baseline estimates from a behavioral model of healthcare spending developed and estimated in Einav, Finkelstein, and Schrimpf (2015). In the figure, we use those estimates to simulate healthcare spending under the standard part D coverage contract shown in Figure 1, and under two counterfactual contracts: a "filled gap" contract (Panel A) and a "no deductible" contract (Panel B). The filled gap contract eliminates the gap by providing pre-gap coverage through the catastrophic limit, while the no-deductible contract eliminates the deductible but offers higher cost sharing (up to the gap) of 38.9 (rather than 25) cents on the dollar, thus leading to the same expected cost to the insurer. To construct the figure, we use the model and baseline estimates from Einav, Finkelstein, and Schrimpf (2015), and simulate 43,000 individuals and associated sequences of health shocks (that are held fixed across contracts). The table entries show the proportion of individuals whose spending ventile gets reshuffled as they move from the standard contract to the alternative contract. Absent heterogeneity in the response to price, all individuals would have stayed in the same spending ventile, and all entries in column (5) would be equal to 1, and all other entries – columns (3), (4), (6), and (7) – would be zero. Heterogeneity in the response to price leads to some individuals moving up and others moving down in the spending distribution in response to the change in coverage.

[a] The first two spending bins are different in size because of the mass point (of 12 percent) of individuals with zero spending under the standard contract.

Appendix Table A2: Illustrative example

	Public Coverage	Private Coverage
1. Individual medical spending	$\lambda_i + \omega_i$	λ_i
2. Individual optimized utilty (u_i*)	$y_i + 0.5 \cdot \omega_i$	y_i
3. Government spending (g_i)	$\lambda_i + \omega_i$	$\lambda_i + \omega_i + s_i$
4. Profits (π_i)	N/A	$\omega_i + s_i$
5. Total Surplus (TS_i)	$y_i - (1+k)\lambda_i - (0.5+k)\omega_i$	$y_i - (1+k)\lambda_i - k\omega_i - ks_i$

Table summarizes the assumptions of our stylized setting in which a private monopoly provider offers a plan that competes with the public default plan. We also assume risk scores are given by $r_i = \lambda_i + \omega_i$ and that the government sets reimbursement for private firms $g_i = r_i + s(r_i)$. That is, the reimbursement is set to be equal to what the government would pay if the individual was covered by the public plan, plus a subsidy s.