Insurance Choice with Non-Monetary Plan Attributes:

Limited Consideration in Medicare Part D

Maura Coughlin^{*} Department of Economics, Rice University Houston, TX 77005, USA

January 2023

Abstract

I propose a model of demand for insurance plans where non-monetary plan attributes stochastically determine the composition of the set of plans that an individual considers, and monetary plan attributes determine the individual's expected utility over contracts in their consideration set. This model reconciles the classic view of insurance contracts as lotteries with purely monetary outcomes with the empirical finding that choice among plans is also driven by other plan attributes. I estimate the model using Medicare Part D data, allowing for unobserved heterogeneity in risk aversion and consideration sets. I find that the latter plays a crucial role in plan choices: while 46 plans are available, more than 90% of individuals consider no more than 5 plans. While the majority of available plans include a deductible, approximately two thirds of all plans considered have no deductible. In contrast to previous literature, I uncover an important role for risk aversion.

Keywords: Insurance demand, limited consideration, Medicare, prescription drugs

^{*}Corresponding author t: (847) 404-3697; e: maura.coughlin@rice.edu

1 Introduction

Choice among products provides individuals with the opportunity to select an alternative that best suits their preferences. In government regulated insurance markets, choice among privately provided plans is also meant to foster competition encouraging higher quality plans and lower costs. In many empirical settings, however, researchers have found that consumers do not select from their choice sets optimally. In some cases, such as health insurance and household electricity plans, researchers have documented evidence suggesting that consumers fail to evaluate all available products at the time of selection and consequently forego available savings (Bhargava, Loewenstein, and Sydnor (2017), Hortaçsu, Madanizadeh, and Puller (2017)). Such choice frictions may also preclude any competitive gains and potentially even give rise to new avenues for firms to increase profits. In some cases, these empirical choice patterns confront the researcher with the gap between decision utility and experienced utility. Of course, the consumer decision making process is generally unobserved and in some instances likely impacted by behavioral biases, the use of heuristics, and unobserved constraints, such as budget or liquidity constraints.¹ Such unobserved factors mean that even if utility is correctly specified, a model without these additional choice process features can easily fail to rationalize the observed choices of individuals. Moreover, in some empirical settings, such as insurance, economic theory provides a specific foundation of demand. At times, empirical choice patterns suggest factors beyond those underlying the theory of demand impact consumers' choices. This departure of individual choice from economic theory can create a modeling puzzle for researchers.

This paper explores this issue in the context of demand for prescription drug insurance, although the overarching motivation and approach are broadly applicable. Health insurance markets in the United States have shifted substantially towards increased consumer choice. Many employers offer their employees a choice from multiple sponsored health insurance plans and expansions and reforms to the Medicare and Medicaid programs, as well as the introduction of ACA exchanges, have expanded the role of choice in publicly regulated insurance markets. The insurance products available and the corresponding choices individuals make in these markets have a large impact on their access to quality healthcare and overall well-being.² An understanding of the

¹See, for example in the case of insurance, Ericson and Starc (2012) and Ericson and Sydnor (2018).

²Prescription drug insurance alone has been shown to improve health outcomes. Diebold (2016) and Semilla, Chen, and Dall (2015), document substantial improvements in drug adherence and mortality rates among beneficiaries enrolled in Medicare Part D.

foundations of individual choice in health insurance markets is crucial to assessing the impact of any new policies, interventions, or modifications to market design, including efforts to improve choice quality.

Rationalizing observed health insurance choices is, however, notoriously difficult. It is not uncommon for individuals to select insurance plans that are strictly dominated by available alternatives.³ The classic economic approach to insurance views contracts as lotteries with purely financial outcomes. Insurance appeals to risk averse individuals as a means to transfer wealth from good states of the world, in which they are not sick, to bad states of the world, where health needs are costly. Numerous studies of prescription drug coverage choice in Medicare have found that beneficiaries appear to overweight premiums relative to out-of-pocket costs and ascribe value to both nonmonetary attributes and monetary attributes above and beyond their out-of-pocket financial impact.^{4,5} It is important to note the challenging nature of this specific environment. During the early years of Medicare Part D, the average beneficiary faced a choice from approximately 50 plans. In a market setting with such a large choice set of complex products, beneficiaries and policymakers alike have expressed concern that the choice environment is difficult to navigate.⁶ There are many potential underlying sources of choice frictions in the Medicare Part D market.⁷ Firm advertising, agent steering, or individual perceptions of firm quality may lead beneficiaries to consider only plans offered by certain firms. Others may simply face cognitive or time limitations that manifest in a reduced number of plans evaluated at the time of enrollment.

In this paper, I propose a model of demand for prescription drug insurance plans where non-monetary plan attributes stochastically determine the composition of the

 $^{^{3}}$ In a relatively simple comparison of employer-provided health insurance where plans differed in deductible and premiums, and thus require a basic dollar comparison, Bhargava et al. (2017) finds a substantial portion of individuals select plans that are strictly dominated regardless of preferences or health realizations. Handel (2013) documents substantial inertia in employer-provided health insurances leading to dominated choices, albeit with reduced adverse selection.

⁴See, for example, Abaluck and Gruber (2011), Ketcham, Lucarelli, Miravete, and Roebuck (2012), and Heiss, Leive, McFadden, and Winter (2013).

⁵I use "non-monetary" in reference to both attributes that do not have an immediate monetary interpretation, as well as to the role of financial attributes above and beyond out-of-pocket costs.

⁶See The Kaiser Family Foundation and the Harvard School of Public Health (2006).

⁷Individuals may face constraints unobserved to researchers that result in the exclusion of certain plans deemed unfeasible or disqualifyingly unappealing. Many individuals face liquidity constraints and are unable to cover large unplanned expenses. See discussion in Durante and Chen (2019) within the section *Dealing with Unexpected Expenses*. It is certainly imaginable that such a constrained individual might only consider plans with reduced or eliminated deductibles. Similarly, many elderly beneficiaries live on a fixed income and a budget-constrained individual may only consider plans with monthly premiums below a reservation price.

set of plans that an individual considers, and monetary plan attributes determine the individual's expected utility over contracts in their consideration set. This model reconciles the classic view of insurance contracts as lotteries with purely monetary outcomes with the empirical finding that choice among insurance plans is also driven by their non-monetary attributes and financial attributes beyond their impacts on costs. This model of limited consideration, in which individuals are assumed to select their preferred plan from an unobserved subset of the feasible set, preserves the structural interpretation of insurance demand as arising from risk aversion, while providing a natural role for various plan attributes to shift choice frequencies in ways beyond the impact of those attributes on the utility derived from a plan. In what follows, "choice set" denotes the full available menu of plans, and "consideration set" refers to the subset of plans an individual compares in terms of expected utility. The model measures the impact of the determinants of limited consideration, such as specific plan attributes, but does not presume a specific underlying behavioral model of consideration set formation.

I estimate the model using Medicare Part D administrative data from the Centers for Medicare and Medicaid Services (CMS). I first verify that the anaylsis sample exhibits plan choice patterns consistent with the puzzles documented in the previous literature. On the whole, individuals are leaving substantial amounts of money on the table by selecting plans that are higher cost for their drug needs, on the order of a quarter of annual spending, and these higher costs are largely not offset by any reduction in variance of expenditure. Individuals place significant weight on plan attributes above and beyond their impacts on the distribution of costs. At times this results in implausible willingess to pay estimates for plan features using standard methods. For example, a full consideration discrete choice model suggests a beneficiary is willing to pay well over a dollar to reduce a plan's deductible by a dollar, even after already accounting for the effect of the deductible on out-of-pocket expenditures.

After documenting choice patterns inconsistent with the standard theory of insurance demand, I estimate the limited consideration model and recover estimates of risk preferences while allowing the probability a plan is considered to depend on plan attributes. A model of expected utility with limited consideration is well suited to explain plan choice patterns among Medicare Part D beneficiaries, and consideration sets play a crucial role in rationalizing plan choices. Beneficiaries in my sample face the choice of 46 plans, but over 90% of individuals consider no more than 5 plans. The probability a given plan is considered is driven by the identity of the insuring firm, the premium, the deductible, and the presence of supplemental coverage in the infamous "donut hole" (a phase of coverage in which beneficiaries pay 100% of drug costs). I estimate the highest premium plan is considered 11% as often as the lowest premium plan, all else equal. Similarly the highest deductible plan is considered 17% as often as a comparable zero deductible plan. The consideration impacts of the firm and deductible alone are appreciable. Three firms account for over 60% of considered plans, while the three smallest account for approximately 0.5%, despite offering a similar number of plans. Although the majority of plans offered in the market include a deductible, around two thirds of considered plans have no deductible. In contrast to the previous Medicare Part D literature, I recover estimates of risk aversion in line with the literature that estimates risk aversion in field data in other insurance markets.⁸ My estimates more than double the mean risk aversion implied by a classic model of full consideration.

The material role of limited consideration, taken together with the distribution of risk aversion, translates into an important cost of limited consideration. Beneficiaries forgo, on average, \$223 in certainty equivalent terms, from considering a subset of plans that often does not include the plan best suited to their drug needs and preferences alone. Of course, it is plausible that beneficiaries are choosing optimally under unobserved constraints. In such a case, this foregone certainty equivalent is better interpreted as a shadow cost of those constraints rather than a welfare loss per se.

This setting, in which economic theory suggests monetary attributes are the only utility-relevant plan features, but empirical patterns are in conflict with that modeling assumption, creates a dilemma for researchers. The model of limited consideration resolves some of the inconsistencies that have become commonplace in modeling insurance choices. This paper provides a intuitive and practical alternative to modeling insurance decisions that both preserves the role of risk preferences and guides policymakers towards how beneficiaries are navigating this complex choice environment. Moreover, the model is tractable to implement even in cases where the choice set is large. My results indicate that documented sub-optimal choice patterns are not a sign that the trend of increasing the role of consumer choice in health insurance is a lost cause but may reflect more nuanced decision-making.

⁸For example, Cohen and Einav (2007) finds among Israeli auto insurance customers a relatively low average risk aversion but a substantial fraction of customers exhibit very high risk aversion. Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2013) finds overall substantial levels of risk aversion among North American auto and home insurance customers.

2 Literature Review

This paper shares a core motivation with previous studies on Medicare Part D plan choice: to understand and evaluate plan choices according to economic models of decision making. Well known studies include Heiss et al. (2013) and Abaluck and Gruber (2011), as well as the exchange resulting from the latter in Ketcham, Kuminoff, and Powers (2016) and Abaluck and Gruber (2016a).⁹ This paper differs methodologically from such prior studies. Both Abaluck and Gruber (2011) and Heiss et al. (2013), among others, evaluate plan choice using a logit approximation to a CARA expected utility function. The latter finds that a specification that includes only variables describing the distribution of out-of-pocket (OOP) costs (mean and variance) fails to explain choice patterns. Moreover, the implied risk preferences are surprisingly unstable over time, with one year of modest risk aversion and one year of substantial risk preference. The former includes other plan attributes to the utility approximation and finds that beneficiaries are selecting plans in a manner inconsistent with expected utility maximization. Specifically, beneficiaries overweight premiums relative to outof-pocket costs, place little to no value on a plan's risk reduction features, and value financial aspects of plans, such as deductible and gap coverage, beyond the impact of such attributes on expected costs. Although incorporating plan attributes into the utility framework improves the explanatory power of the model, the resulting estimates are challenging to interpret in the classic insurance model.¹⁰ Inclusion of such attributes such can offer insight into which plan features relate to choice probabilities but at times at the expense of a straightforward economic interpretation. The expected utility model with limited consideration proposed in this paper resolves this tension.

Given the complexity of prescription drug plans, the large number of plans available, and the possibility of constraints impacting the beneficiaries, it is hardly surprising that individuals would fail to behave in a manner fully consistent with standard economic models. The literature on menu complexity and heuristic shortcuts in insurance also relates to this paper. In the Part D market, Ketcham, Lucarelli, and Powers (2015)

⁹Additionally important early studies include Heiss, McFadden, and Winter (2010), Lucarelli, Prince, and Simon (2012), Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2012), Ketcham et al. (2012), and Kesternich, Heiss, McFadden, and Winter (2013), among others.

¹⁰There are reasons to assume a beneficiary may ascribe a "cost" to certain attributes - for example, paying a deductible may cause disutility due to liquidity constraints that make a large single payment particularly challenging. Adding the deductible as a term in the utility specification, however, suggests a constant utility "cost" of the attribute across all possible health realizations, similar to the premium. See Handel (2013) and Handel and Kolstad (2015) for a discussion of this topic.

finds evidence that it is not the size of the choice set alone that drives choices inconsistencies. The quality of choices is estimated to improve with larger choice sets due to increased switching, with the exception of cases where additional plans are relatively more expensive. In other health insurance markets, there is evidence that consumers use heuristic shortcuts to limit the choice set before choosing plans (Ericson and Starc (2012)), as well as that choices improve when products are standardized and the choice set becomes less complex (Ericson and Starc (2016)).

Inconsistencies with model implications can also be suggestive of model misspecification. Ketcham et al. (2016) implements a very general test of rationality, using General Axiom of Revealed Preference (GARP) arguments to determine if Part D plan choices are consistent with any utility specification. However, even under such a general framework a sizable fraction of initial plan choices remain inconsistent with utility maximization.¹¹ Such deviations from rationality are not unique to prescription drug insurance choices. Bhargava et al. (2017) describes a case of employer offered health insurance plans in which a substantial portion of individuals select insurance plans that are strictly dominated by available alternatives, and in such an unambiguous manner that basic arithmetic would highlight that dominance.¹² In the market for auto collision insurance Barseghyan, Molinari, and Thirkettle (2021b) and Barseghyan, Coughlin, Molinari, and Teitelbaum (2021a) document a substantial fraction of individuals selecting a policy that is dominated by other available plans, regardless of risk preferences.

A commonly suggested and intuitive explanation for the prevalence of what economists deem suboptimal choices is limited consideration.¹³ In a model of limited consideration, individuals are assumed to select a plan (or product, more generally), from a considered subset of the feasible set. Choices, therefore, do not reveal preference over the entire choice set, but rather only over the considered set. Previous studies, including Abaluck and Gruber (2011) and Abaluck and Gruber (2016b), have mentioned limited consideration as a possible explanation for the role of plan attributes in choices.¹⁴ A

 $^{^{11}}$ In the main version of this test, 21% of choices over 2006-2010 were inconsistent with utility maximization. The further relaxed consistency test found 14% of those choices are not rationalizable.

¹²The setting included pairs of plans differing only in the deductible, but the additional premium charged for the lower deductible plans exceeded the amount by which the deductible was reduced, guaranteeing larger costs under all realized health scenarios.

¹³Models of limited consideration and limited attention have a long history in economics, including Simon (1959).

¹⁴Both studies posit that the importance of firm fixed effects in matching choice patterns may suggest individuals are rationally using trusted firms as a heuristic shortcut when unable or unwilling to make the time-consuming or difficult financial comparison across all plans.

well-studied form of limited consideration in the Medicare Part D market is inertia, a pattern of behavior in insurance markets whereby individuals passively remain in existing plans at the time of a renewal and a model of limited consideration in which the consideration set is either empty or contains the singleton default. Ho, Hogan, and Scott Morton (2017), studies Medicare Part D choices over time and documents the role of inertia and the way in which certain shocks - most notably in premium, a highly visible plan attribute - can break beneficiaries from their inertia. Abaluck and Gruber (2016b) finds a role for inertia and little evidence of learning or improved performance as beneficiaries gain experience. Polyakova (2016) explores the interaction of inertia, adverse selection, and market regulations and finds inertia and switching costs contribute to the sustainment of an adversely-selected equilibrium in Medicare Part D.

Keane, Ketcham, Kuminoff, and Neal (2019) and Ketcham, Kuminoff, and Powers (2019) propose an alternative approach in which Part D choices are assumed to be made with varying degrees of consumer informativeness. In the former, a mixture-of-experts model is used to model plan choices as probabilistically revealing of preferences. The latter uses survey data as a signal of whether consumers are informed and assesses the welfare implications of various market interventions assuming the observed choices of informed individuals proxy for the preferences of uninformed individuals. In contrast, Brown and Jeon (2020) build on the work of Matejka and McKay (2015) and Fosgerau, Melo, Palma, and Shum (2017) and propose a model of rational inattention whereby beneficiaries for whom the choice stakes are high, those with high variance of out-ofpocket costs across available plans, acquire more information about the plans before enrollment. A related study of limited consideration in Abaluck and Adams-Prassl (2021) focuses on identification of a default specific model of limited consideration and an alternative specific random consideration model arising from demand asymmetries. The authors estimate a hybrid model in the setting of Medicare Part D and find that choices in the market are consistent with their model of limited consideration.

Beyond health insurance, this paper builds on the methodology of limited consideration in discrete choice models. The alternative specific consideration model used below has been developed and shown to be nonparametrically identified under certain conditions in Barseghyan et al. (2021b).^{15,16} This paper highlights a major appeal of

 $^{^{15}}$ The model in Barseghyan et al. (2021b) expands on the work of Manski (1977) and Manzini and Mariotti (2014).

 $^{^{16}}$ A conceptually similar, but econometrically different and parametrically specified, model was used in Goeree (2008) to estimate demand for computers when advertising plays a role in consideration set formation.

such models to empirical applications. The incorporation of consideration sets provides a manner for researchers to model product choice in cases where both preferences and context-specific forces are at play by introducing a natural role for variables that are outside the underlying economic model of utility. This is the key methodological innovation of this paper compared to previous studies. In such an insurance model, there is a distinction between what enters expected utility and reflects the uncertainty of the environment and the plan attributes that enter consideration and hence do not depend on the state of the world. In cases where plan attributes impact choice beyond utility-relevant monetary costs, consideration sets provide a theoretically sound avenue to relate attributes to choice. A limited consideration model can marry the theoretical underpinnings of expected utility with the empirical reality present in this market. These techniques offer a tractable modeling alternative, even in the presence of such a large feasible choice set.

3 Model of Prescription Drug Insurance Choice

In the empirical application below, each individual has uncertain health (and therefore prescription drug) needs during the coverage year. The coverage and costsharing structure of each available prescription drug plan translates those drug needs into OOP costs. Plans differ in whether and how generously drugs are covered (by classifying drugs into a coverage tier), how cost-sharing is applied to different tiers of drugs, and at which levels of costs an individual moves between different phases of plan coverage. Motivated by empirical findings in previous studies noted in Section 2 and evidence in the present study's sample provided below in Section 4.3 that numerous plan attributes beyond OOP drug costs affect individuals' choices, this paper proposes an expected utility model with limited consideration to disentangle the role of risk preferences and plan attributes in the demand for prescription drug insurance.

My model of insurance plan choice adopts the expected utility framework standard in the literature on decision-making under uncertainty, although the limited consideration model can be applied to non-expected utility models as well. Individuals $i \in \{1, 2, ..., I\}$ face a choice of plans indexed by $j = \{1, 2, ..., J\}$. Individuals have utility over final wealth represented by utility function U(x) and face distributions of financial losses that vary across plans j. Plans are differentiated by their premiums p_j and plan benefit design, which translates drug claims into OOP costs. Each plan includes a list of drugs, known as the plan forumlary, with a classification of coverage generosity (tier) and a negotiated list price of the drug to which costsharing is applied. Each plan also designates a copay or coinsurance amount for each tier of covered drugs. Plans also differ in their designation of coverage phases across which different costsharing rules apply. Taken together, plan design converts a sequence of drug claims, along with premiums, into a total yearly OOP cost C_{ij} , and thus the same drug needs will differ in costs across plans. Each individual *i* faces a distribution of costs under each available plan, $F(C_{ij}; \theta_{ij})$, knowable from data. Denoting initial wealth by W_i , the realized, year-end wealth of an individual is given by $W_i - C_{ij}$. Individuals are heterogeneous in their taste for risk and are characterized by a coefficient of risk aversion, ν_i . Individuals compare insurance plans on the basis of expected utility.

ASSUMPTION 3.1 (Plan Costs and Utility):

- 1. Plans are characterized by cost-sharing attributes \mathbf{Z}_{j} .
- 2. The out-of-pocket costs an individual faces in each plan during the year is a function of their uncertain health needs and the cost attributes of the plan, $C_{ij} = f(health_i, \mathbf{X_j})$. $F(C_{ij}; \theta_{ij})$ is recoverable from data.
- 3. Individuals are expected utility maximizers with $U_i(x)$ drawn from a known set of utility functions $U_i(C_{ij}) = U(W_i C_{ij}; \nu_i), \ \nu_i \sim F_{\nu}$.
- 4. Individuals are expected utility maximizers and evaluate the plans they consider according to expected utility:

$$EU_i(C_{ij}) = \int U(W_i - C_{ij}; \nu_i) dF(C_{ij}; \theta_{ij}).$$
(3.1)

Individuals enroll in the plan that delivers the greatest EU_i .

The standard model assumes that each individual compares all available plans and selects the plan that delivers the highest expected utility. As described in Section 2, previous studies have found that plan attributes determine choice in ways beyond the experienced financial impacts of those attributes but have struggled with a rational structural utility explanation for such a role. This model posits that these important, but not directly utility-relevant, variables impact choice by determining the composition of the consideration set an individual evaluates when selecting a plan. Such a model preserves the theoretical structural underpinning of insurance demand through an expected utility framework while reconciling the empirical reality that risk preferences alone cannot rationalize insurance choices in some markets. Individuals are assumed to select an unobserved subset of the feasible choice set to actively consider, denoted $M_i \subseteq \mathcal{M}$, and select a plan from that subset.

Plan attributes are modeled to impact utility only insofar as they impact the distribution of monetary costs individuals face under each plan. For example, consider the simplest case in which an individual facing the prospect of a loss $L \sim F(L)$ can choose between two otherwise identical insurance plans with no deductible and constant coinsurance rates, $c_1 > c_2$, and corresponding premiums $p_1 < p_2$. Expected costs are $p_1 + c_1 \int LdF(L)$ under the first plan and $p_2 + c_2 \int LdF(L)$ under the second plan. The premiums and coinsurance rates are assumed to not influence expected utility beyond their impacts on the distribution of monetary costs. In other words, individuals do not derive utility from low premiums specifically, for example, beyond the direct impact of a low premium on costs.

In many settings, however, monetary costs alone do not explain observed choice patterns well, and researchers have found other non-monetary attributes to impact choice probabilities. I use the term "non-monetary" to refer to all additional product attributes other than C_{ij} . This includes product attributes that do not have an immediate monetary interpretation, such as the name of the product or firm offering the product, as well as monetary attributes above and beyond their impact on the distribution of monetary costs, such as a deductible amount in and of itself. In this model, such plan attributes, denoted \mathbf{Z}_{j} , enter the decision process through the formation of consideration sets M. Specification of the consideration set formation process can be tricky without additional data or an experimental setting. Rather than model a specific consideration set formation process, I use an alternative specific consideration probability model, similar to the models found in Barseghyan et al. (2021b) and Manzini and Mariotti (2014), among others, as a reduced form estimate of consideration. Each plan appears in an individual's consideration set with a plan-specific consideration probability $\varphi_i(\mathbf{Z}_i)$. Product consideration probabilities are homogeneous across agents facing the same feasible choice set.¹⁷ Conditional on observables, each product's appearance in a consideration set is assumed to be independent.

ASSUMPTION 3.2 (Limited Consideration):

 $^{^{17}}$ This assumption can be relaxed by allowing φ_j to depend on individual attributes, such as age or sex, given sufficient data.

1. Each individual i selects a subset of the available plan menu, $M_i \subseteq \mathcal{M}$, evaluates the expected utility of the plans, and selects their utility preferred plan:

$$j^* = \underset{j \in M_i}{\operatorname{arg\,max}} EU_i(C_{ij}). \tag{3.2}$$

2. Consideration set formation follows an alternative specific random consideration model. Each plan j appears in M_i with probability $\varphi_j(\mathbf{Z_j})$, independent across plans conditional on attributes and homogeneous across individuals. Thus, the probability of any given consideration set is:

$$Pr(M_i = M) = \prod_{k \in M} \varphi_k(\mathbf{Z}_k) \prod_{k' \notin M} \left(1 - \varphi_{k'}(\mathbf{Z}_{k'})\right).$$
(3.3)

- 3. Consideration is assumed independent of risk aversion.
- 4. In the case $M_i = \emptyset$ a completion rule is assumed.

The probability *i* selects *j* is jointly determined by the likelihood *i* drew a consideration set that contains plan *j* and that among the plans in M_i , *j* was the most preferred. Denoting *i*'s choice of plan j^* by $y_{ij^*} = 1$, this choice probability can be written as (suppressing conditioning notation for simplicity):

$$Pr(y_{ij^*} = 1) = \sum_{M \subseteq \mathcal{M}: j^* \in M} Pr(M_i = M) Pr(EU_i(C_{ij^*}) \ge EU_i(C_{ik}) \forall k \in M)$$
(3.4)
$$= \sum_{M \subseteq \mathcal{M}: j^* \in M} \prod_{k \in M} \varphi_k(\mathbf{Z}_k) \prod_{k' \notin M} (1 - \varphi_{k'}(\mathbf{Z}_{k'})) Pr(EU_i(C_{ij^*}) \ge EU_i(C_{ik}) \forall k \in M)$$
(3.5)

For simplicity the completion rule component of probability is left implicit, but in the event an individual draws $M_i = \emptyset$, in the empirical application, I assume one of the available plans $j \in \mathcal{M}$ is drawn as the consideration set with equal probability.

While the distribution of C_{ij} captures the utility-relevant features of each plan, plan characteristics \mathbf{Z}_{j} include additional attributes:

$$\mathbf{Z}_{j} = \{\operatorname{firm}_{j}, \operatorname{premium}_{j}, \operatorname{deductible}_{j}, \operatorname{gap}_{j}, \operatorname{Top100}_{j}, \operatorname{AvgCS}_{j}\},\$$

where gap_j is an indicator for supplemental coverage in a market-specific high spend coverage phase, $Top100_j$ is an index of formulary generosity, and $AvgCS_j$ is an index summarizing the various copay and coinsurance rates included in plan j's design. These attributes are discussed in detail in Section 4.1. Some elements of \mathbf{Z}_{j} are related to C_{ij} , such as premium_j. Other elements have a complex, nonlinear relationship, such as deductible_j, gap_j, Top100_j, and AvgCS_j, where the relationship depends heavily on the individual-specific drug needs and the plan formulary.¹⁸ Note firm_j does not determine C_{ij} . This paper leverages results existing in the literature about the use of exclusion restrictions between determinants of utility and determinants of consideration in such models. Identification is discussed below in Section 5.4.

For intuition on how plan attributes determine consideration, it is helpful to consider only two attributes: deductible and premium. An individual may limit their consideration of plans based on these attributes for a number of plausible reasons. In the presence of liquidity constraints, an individual may be unable to afford large lump expenses and only consider plans with low or zero deductibles. If an individual faces budget constraints, they will not consider plans with monthly premiums in excess of a reservation price. More generally, individuals are modeled as simply less likely to consider a plan with less desirable attributes than one with better attributes. There is a clear ordering of deductible and premium: lower is unambiguously better than higher, all else equal. Figure 3.1 visually demonstrates the connection between the desirability of a plan attribute and consideration probabilities. In this example, deductible and premium each take one of three values: low, medium, or high. The darker bottom left region corresponds to the best plans along these two attributes - those with the most preferred low deductible and the most preferred low premium. For expositional simplicity, plan consideration probabilities are presented as the product of the attribute specific probabilities. Consideration is modeled to diminish as plans move further away from best along each attribute dimension. The lightest shaded box in the upper right corner corresponds to plans with both the highest premium and highest deductible and are, thus, least likely to appear in an individual's consideration set.

With the exception of the insuring firm, all of the included attributes have such an ordering. All else equal, lower deductibles, premiums, and cost-sharing is preferred. Similarly covering more drugs is preferred to fewer, and supplemental gap coverage is better than none. The intuition of Figure 3.1 is applied across these multiple dimen-

¹⁸Many plans with deductibles do not apply them uniformly - some drugs are exempt from deductible and have OOP costs immediately according to the cost-sharing of the initial coverage phase. Thus, even this simple attribute enters OOP in a complicated, individual-specific way.



Figure 3.1: Consideration Intuition with 2 Attributes

Note: $p_1 < p_2 < p_3$ represents the consideration probabilities for each deductible and $q_1 < q_2 < q_3$ the consideration probabilities for premiums.

sions. In the absence of an objective ranking over firms, and to reflect the numerous underlying mechanisms causing individuals to consider firms differentially, the model assumes a base consideration probability for each of the firms in the market. It is to this base probability that the reductions in consideration according to attributes is applied. The details of the parameterization are discussed in Section 5.2.

4 Application: Medicare Part D

4.1 Institutional Background

Since 1965, Medicare provides hospital (Part A) and physician services (Part B) insurance coverage for elderly Americans and those with disabilities and certain serious illnesses. In 2006, as part of the Medicare Modernization Act of 2003, prescription drug coverage was added to the program. Beneficiaries seeking prescription drug coverage have the option of enrolling in a standalone prescription drug plan (PDP) through Medicare Part D or to bundle prescription coverage with the other health insurance through Medicare Advantage (also known as Medicare Part C). Both Medicare Advantage and Part D are regulated by CMS but provide beneficiaries choice among plans offered by private insurance companies.

Participants in Part D select plans annually during open enrollment. Generally

those who do not qualify for low-income subsidies cannot change their plan throughout the year. The menu of available plans is determined based on in which CMS region a beneficiary resides. Within each region, beneficiaries face a large number of plans to choose from, with premiums subsidized by the federal government and fixed across individuals. In 2010, regional choice sets varied from a minimum of 39 plans to a maximum of 54 plans. Firms participating in a given market can offer multiple plans and have some discretion over ways to differentiate their plans.

All plans offered through the program must meet CMS requirements on minimum plan generosity, including covering at least 2 drugs within 148 therapeutic categories, and virtually all drugs within certain crucial therapeutic classes. Every year CMS releases cost-sharing standards for a base plan design. Plans are required to be at least as generous actuarially as the standard plan. The standard plan divides beneficiary spending into four phases: the deductible, the initial coverage phase, the coverage gap (known colloquially as the "donut hole" and phased out as part of the Affordable Care Act), and the catastrophic coverage phase. Figure 4.1 provides a graphical representation of the 2010 standard plan. There are multiple ways a firm can differentiate the plans they offer from one another and from those offered by other firms. Insurers have discretion over the formulary, which lists all drugs covered under a plan and how generously they are covered by classifying each included drug into a tier (lower tiers correspond to lower cost drugs). Firms can also adjust the cost-sharing structure of a plan, with many choosing to use copays over coinsurance rates, and to offer plans with a reduced or fully eliminated deductible.

Despite the complexity of the market, Part D has been, on the whole, lauded as a success. Studies, including Diebold (2016) and Semilla et al. (2015), have found substantial improvements in prescription drug adherence and mortality rates among beneficiaries enrolled in the program. The program is popular among beneficiaries, with 49 million beneficiaries enrolled in 2022. From the program's inception, however, there has been concern that the plan choice environment is overly complex, especially for the target population. Beneficiaries themselves expressed interest in a reduced choice set in order to alleviate the difficulty or time involved in choosing a plan.¹⁹ The number of plans offered has decreased from the initial years of the program with the

¹⁹For example, The Kaiser Family Foundation and the Harvard School of Public Health (2006) notes that in the first year of the program 73% of seniors found the program too complicated, as did 91% of pharmacists and 92% of doctors surveyed. 60% of seniors agreed that Medicare should select a small number of good plans to help seniors have an easier time choosing.



Figure 4.1: 2010 Standard Plan Design

Notes: Any claims beyond the out-of-pocket threshold (\$4,550) are treated as catastrophic and the beneficiary pays the maximum of 5% coinsurance or a copay of \$6.30 for branded drugs and \$2.50 for generic drugs.

average beneficiary now facing a set of approximately 30 plans.

4.2 Data

The primary data source is administrative data from CMS. These data include information for a 5% random sample of 2010 Medicare beneficiaries. The relevant enrollee data include information on basic demographics, plan choice, and the full set of drug claims filed under the beneficiaries' plans. These beneficiary and claims data are paired with plan information, linking premiums and plan coverage structures for all plans available to each beneficiary. Additional information on the formularies and base drug prices is included in public use files released for purchase by CMS.

4.2.1 Analysis Sample

This study requires restricting the sample to beneficiaries selecting standalone prescription drug plans (PDPs), excluding those who forgo prescription drug coverage, those with coverage outside of Medicare and those who opt instead for Medicare Advantage plans. Additionally, I exclude all individuals receiving a low-income subsidy. The enrollment, pricing, and choice environments for those individuals differ substantially from the standard Medicare population, and I lack relevant data. I also exclude those with more than one Medicare drug plan over the year, are dual eligible for Medicaid, or drop their coverage mid-year for any reason other than death. As a final restriction, I exclude individuals who either currently have or initially enrolled in Medicare due to end-stage renal disease.

Every annual open enrollment, beneficiaries select a plan for the entirety of the following year. If a beneficiary's existing plan remains available, they default into the same plan unless they make an active choice to switch. A concern in any choice environment of this sort is the distinction between the role of preferences and the role of inertia in observed choices. I abstract from concerns of inertia by restricting attention to beneficiaries making "active choices." Active choices include enrollment decisions of those joining Medicare Part D upon eligibility and those first choosing a Part D plan but for a few common reasons, not making that choice at the time of eligibility. This includes those that either retained employer drug coverage for a period of time after entering Medicare, initially retained other creditable prescription coverage while Medicare enrolled, or went for a period of time without any drug coverage. If these individuals joined Part D during 2010, their choices are included in my sample. The final group included are those who are actively switching plans from the previous year.²⁰ It is worth noting that, as shown in summary statistics, these sample selection criteria unsurprisingly result in a sample that is younger than the overall Medicare population. To the extent that the model presented here captures behaviors one might expect to be more pronounced among the general Medicare population, this analysis may underestimate the role of limited consideration.

The set of plans available is determined by the region of residence. This study focuses on active choices among residents of California. Table 4.1 presents summary statistics of the full 2010 active choosers sample and the 2010 California subsample. Californians in 2010 could choose from 46 plans, offered by 19 different firms. On average, California beneficiaries are less white, slightly younger, and, correspondingly, file fewer claims. In a similar fashion to the national average, these beneficiaries are largely choosing plans offered by the most popular firms. These individuals, however, enrolled in zero deductible plans more often than the average active chooser in the

 $^{^{20}}$ Without the Part D Plan Election Type Beneficiary Summary File, it is difficult to determine if a plan change is an active choice or a passive transition upon termination of the existing plan. Conservatively, I include as active choosers those switching plan types, for example from an HMO to a PDP, from 2009 to 2010, as well as those who select a 2010 plan offered by a different firm than their 2009 plan.

US. The majority of the 46 plans available to the California beneficiaries include a deductible, yet 69% of the sample enrolled in a plan without a deductible.

v		
	U.S.	CA
Sample Size	69,217	4,513
2010 Months Covered	9.9	9.4
Age	71.0	69.9
Female	.584	.565
White	.937	.890
Monthly Claims	2.5	2.3
Average Total OOP	\$618	\$628
Number Plans Offered	46.6	46
Avg Deductible	\$96.76	\$65.58
Zero Deductible	.534	.678
Avg Monthly Premium	\$35.30	\$37.67
Top 1 Most Popular Firm	.320	.367
Top 2 Most Popular Firms	.509	.512
Top 3 Most Popular Firms	.618	.641

Table 4.1: Summary Statistics: Active Choosers

Note: Statistics computed over "active choosers" in the 2010 sample based on description above. All statistics reflect unweighted averages and include beneficiaries who are enrolled for a portion of the calendar year.

The plans offered in the California market exhibit substantial variation in the attributes previously found to be choice-relevant in the literature. The 46 available plans were provided by 19 different insurance firms. 16 plans included the maximum deductible, 11 included a reduced deductible, and 19 did not include a deductible. 20% of the plans offered some form of coverage in the donut hole. Of the 100 most popular drugs by sale among beneficiaries, the plans in California covered between 71 and all of them, with an average of approximately 91 drugs covered. The plans offered varied in average cost-share in the initial coverage phase from 33% to 58%. Among this large and varied choice sets, beneficiary choices were fairly concentrated. Only 16 of the 46 plans garnered enrollment in excess of 1% of the sample.

4.2.2 Distribution of Expected Out-of-Pocket Costs

To estimate the model of plan choice, I require the distributions of counterfactual OOP costs beneficiaries would face under the full set of alternatives available to them. To estimate these counterfactual costs, I construct a plan calculator that takes in any specified set of claims for an individual and computes the OOP expenses that the specified sequence of claim events generates under every plan available. Consider an individual who fills a number of prescriptions each month. Under each available plan's formulary, each of those drugs is classified by tier that determines the cost-sharing structure used, whereby cheaper drugs are assigned a lower tier than more costly drugs. Additionally, in each plan a different base price of the drug has been negotiated to which the plan's cost sharing structure is applied. The calculator procedure involves determining the tier each plan assigns a drug and calculating the out of pocket costs for each claim accounting for the cumulative costs and corresponding coverage phase.

The purpose of the calculator is to quantify counterfactual spending distributions, as an economically rational beneficiary would compare plans in terms of the OOP costs in each plan. It is not obvious, however, what sequence of claims an individual anticipates at the time of plan choice. Some have assumed that beneficiaries have "perfect foresight", and assume that at the time of plan choice, beneficiaries compare the OOP costs of the drugs they would come to claim during the year of coverage.²¹ Alternatively, some studies have assumed a myopic approach, assuming that beneficiaries base their expectations on their previous year drug claims when data is available or current drug needs.²² It is also possible to take a "rational expectations" approach and assume individuals predict their drug needs will be drawn from a distribution of costs under each plan based upon the realized costs of a set of "similar" individuals. For expected mean expenditures, the results in Section 6 adopt a perfect foresight assumption. The robustness analysis in Online Appendix B.1 presents results under a myopic approach by projecting the first month of claims experience in 2010 for the remainder of a beneficiary's time in the plan (note the popular and CMS-promoted online tool to help with plan choice, PlanFinder, uses this approach), an alternative myopic approach by projecting the first two months of 2010 claims experience, and a rational expectations approach based on the binning procedure described in the following paragraph.

Any measure of higher order moments of the distribution of expected costs requires

²¹For example, ex post claims are used as the anticipated mean out-of-pocket costs under plans in Abaluck and Gruber (2011) and as one of two alternative models in Abaluck and Gruber (2016b).

²²See, for example, Kesternich et al. (2013), Heiss et al. (2013), and Abaluck and Gruber (2016b).

additional calculations. To estimate a distribution of OOP costs an individual in the analysis sample expects, beneficiaries are grouped into bins of "similar" individuals based on average monthly number of claims and average monthly gross cost of claims. In cases where claims correspond to multiple months' supply of drugs, I treat it as multiple claims. Details of this procedure are outlined in Online Appendix A. A random sample is drawn from each bin, and their claims are passed through the plan cost calculator to estimate a distribution of costs under each plan. The higher order moments of the cost distribution that enter an individual's utility function are computed from this sample distribution of similar individuals.

4.3 Empirical Motivation for Limited Consideration

In many previous studies of Part D enrollment, even though the specifics of the data and models may differ, there is broad evidence that seniors are selecting drug plans that are more expensive for their drug needs than available alternatives. It is challenging in empirical settings to distinguish between preferences and consumer "mistakes." To motivate the estimation of the model described in Section 3, I conduct reduced form analyses to show that the patterns of choice inconsistencies documented in previous studies manifest in my sample. Table 4.2 presents a GARP-style test of rationalizability following Ketcham et al. (2016). I compute the share of individuals selecting plans on the mean of total OOP (including premium) expenditure frontier, the mean-variance of OOP expenditure frontier, and the mean-variance-firm frontier. By focusing on dominance, these measures test whether choices are consistent with *some* utility function rather than a certain specification. Fewer than 17% of beneficiaries select the lowest cost plan for their realized drug needs. I find approximately 24% of beneficiaries select a plan within 5% of their minimum cost plan, and around 30% within 10% of the minimum cost plan.

It is consistent with standard insurance theory for a risk averse individual to pay more in mean costs for reduced variance in expenditures. Evaluating plan choices on the mean-variance frontier implies choices are dominated only if there is another plan available that is at least as good in terms of mean and variance of expenditures and strictly better in at least one of those measures. Rationalizability of observed choices improves by this measure but the majority of plan choices remain dominated. This means that, on the whole, the foregone savings are not offset by a reduction in risk. A

	% of Sample
Rationalizability Test	
Mean Frontier	17.4
Mean-Variance Frontier	43.6
Mean-Variance-Firm Frontier	88.5

Table 4.2: Choice Rationalizability and Cost Distribution Frontier

Notes: Mean assumes perfect foresight. Variance estimated from a distribution of 100 randomly sampled "similar" individuals as described in Online Appendix A.

major boost to explanatory power comes from incorporating the firm. In this test 89% of plan choices are consistent with utility maximization of some utility function. This test of rationality designates a choice as a mistake if a beneficiary selects a plan that is dominated in the mean-variance space by another plan offered by the *same* firm. In this market, each firm offers typically 1-3 plans within a region, leaving little room for a dominating plan. And yet, 11% of this sample selects such a dominated plan.

A related pattern that emerges in the lower panel of Table 4.1 is the prevalence of certain attributes among chosen plans. The majority of California beneficiaries select a plan without a deductible, even though, as described in Section 4.2.1, the majority of plans offered in California include a deductible. Beneficiaries are on the whole selecting plans with low deductibles, low premiums, and offered by one of the three most popular firms. As a statistical test of explanatory relevance, Table 4.3 presents results of a simple logit regression. Column (1) includes in the regression the monetary variables describing the distribution of costs included in a standard model of insurance demand. Column (3) includes plan attributes $\mathbf{Z}_{\mathbf{j}}$. The Pseudo R^2 of the regression in Column (3) is approximately three times that of Column (1). Columns (2) and (4) allow the coefficients on premium and expected OOP costs to differ. Figure 4.2 graphically contrasts the explanatory power of the logit regressions in Columns (1) and (3) by plotting the implied choice probabilities under each set of estimates. The right-most bar is a composite plan aggregating all 17 plans in which between 1 and 10 individuals in the sample enrolled. The improvement in fit with the additional attributes is visually obvious.

		0		
	(1)	(2)	(3)	(4)
$Premium + \mathbb{E}OOP$	-0.522***	—	-0.389***	_
(hundreds)	(.009)		(.011)	
Premium	_	-0.599***	—	-0.579***
(hundreds)		(.013)		(.027)
EOOP	_	-0.461***	_	-0.353***
(hundreds)		(.011)		(.011)
Variance	-0.031***	033***	-0.002	-0.0002
	(.002)	(.002)	(.002)	(.003)
Deductible	_		-0.630***	-0.681***
(hundreds)			(.026)	(.027)
Gap	_		-0.793***	-0.313***
			(.077)	(.094)
Top100 Drugs	_		-0.065***	051***
			(.006)	(.006)
Avg CS	_		0.080	-0.627
			(.625)	(.624)
Firm Dummies	No	No	Yes	Yes
Pseudo \mathbb{R}^2	0.137	0.139	0.357	0.359

Table 4.3: Conditional Logit Estimates: Impact of Non-Monetary Attributes

Notes: Standard errors are in parentheses. Variance denotes the variance of $\mathbb{E}OOP$ measured in hundreds of dollars. *** Significant at 1% level.



Figure 4.2: Logit Implied Choice Distribution

Notes: Panel (a) plots the model implied choice probabilities from Column (1) of Table 4.3 in red. Panel (b) plots the model implied choice probabilities from Column (3) in red. In both figures, the blue bars correspond to the empirical choice shares. Plans are ordered from the plan with the largest enrollment share on the left to the plans with zero enrollment. The rightmost plan corresponds to a composite plan of the 17 plans in which between 1 and 10 individuals enrolled.

With the exception of the measure of average cost-share, the additional plan attributes in Columns (3) and (4) are highly significant. Some of these attributes are non-monetary attributes. There is no immediate way to compare these variables to those related to the costs of each plan. The deductible, while financial in nature, is also not directly related to costs. Insofar as the deductible of a plan impacts the costs under each plan, it is already accounted for in the expected OOP cost term. The coefficients on the plan attributes reflect a relationship between the attributes and plan choice above and beyond their impacts on OOP costs. These estimates only reflect correlation but are informative for the structural model. If one were to take these estimates as reflecting underlying demand, the estimates are rather challenging to interpret, although a model excluding them explains choices poorly. As discussed further in Online Appendix B.2, these results suggest, for example, that beneficiaries are willing to pay well over \$310 to reduce a \$310 deductible to zero. In addition, in Online Appendix B.3 I present results from basic empirical tests suggestive of limited consideration.

There are many possible mechanisms through which these non-monetary attributes can affect plan choice. Some may find the large menu of plans burdensome and employ heuristics to reduce the choice set to a more manageable size. Others may have liquidity constraints and only consider plans with a reduced or eliminated deductible. Premiums may receive additional weight over expected out-of-pocket costs due to budget constraints. The model here does not require the researcher to take a stance on how exactly beneficiaries are considering plans based on their attributes. This agnostic approach is focused on flexibly approximating this process in order to learn what beneficiaries appear to be weighing when they make Part D plan choices and leaves to future work more precise exploration of the details underlying the consideration set formation process in this market.

5 Estimation

5.1 Expected Utility Specification

As is common in the literature on insurance demand, I assume individuals are risk averse. I further assume that risk aversion is constant across values of wealth and heterogeneous across agents with $\nu_i \sim F(\nu; \theta_{\nu})$. This emits a utility model of constant absolute risk aversion (CARA).

ASSUMPTION 5.1 (Utility Specification):

1. Individuals have CARA utility over final wealth:

$$U_i(C_{ij}) = -\frac{1}{\nu_i} \exp(-\nu_i (W_i - C_{ij})) = -\frac{1}{\nu_i} \exp(-\nu_i W_i) \exp(\nu_i C_{ij}).$$
(5.1)

- 2. Risk and risk preferences are independent of one another.
- 3. $C_{ij} \sim N(\hat{\mu}_{ij}, \sigma_{ij}^2)$.
- 4. $\nu_i \in [0, .01] \sim Beta(\beta_1, \beta_2).$

Rather than incorporating a random error into utility, stochastic choice, conditional on preferences, arises through the formation of the consideration set. The omission of the idiosyncratic component of utility here stems from the critique of such random utility models for choice under risk and uncertainty (see Apesteguia and Ballester (2018)) and additionally simplifies estimation of the model substantially, as described below. Assuming that risk and risk preferences are independent and conditional on a beneficiary's coefficient of absolute risk aversion, expected utility takes the form:

$$EU_i(C_{ij}|\nu_i) = -\frac{1}{\nu_i} \exp(-\nu_i W_i) \mathbb{E}\left(\exp(\nu_i C_{ij})\right).$$
(5.2)

Note that for a fixed value of ν_i , $\mathbb{E}(\exp(\nu_i C_{ij}))$ is the moment generating function of the random variable C_{ij} . Similarly to elsewhere in the literature, OOP costs are assumed to be Normally distributed, $C_{ij} \sim N(\hat{\mu}_{ij}, \sigma_{ij}^2)$, where $\hat{\mu}_{ij} = p_j + \mu_{ij}$ is the mean OOP drug expenditures of individual *i* under plan *j*, shifted by the personinvariant premium for plan *j*. The cost parameters μ_{ij} and σ_{ij}^2 are computed outside of the model as described in Section 4.2.2 and Online Appendix A. Substituting for the moment generating function, and denoting $EU_i(C_{ij})$ with the simpler EU_{ij} , expected utility can be written as a function of the mean and variance of out-of-pocket costs:

$$EU_{ij} = -\frac{1}{\nu_i} \exp(-\nu_i W_i) \exp(\nu_i \hat{\mu}_{ij} + \frac{1}{2} \nu_i^2 \sigma_{ij}^2).$$
(5.3)

Relative utility and the ordinal ranking of plan utilities are not impacted by the positive multiplicative term $\frac{1}{\nu_i} \exp(-\nu_i W_i)$. Thus, this value can be divided away from all utility levels and utility rankings remain unchanged. Therefore, for estimation purposes, a simpler form of expected utility suffices:

$$EU_{ij} = -\exp(\nu_i \hat{\mu}_{ij} + \frac{1}{2}\nu_i^2 \sigma_{ij}^2).$$
 (5.4)

5.2 Consideration Probabilities

As described in Section 3, while the distribution of OOP costs impact expected utility, the model allows plan choice to depend on \mathbf{Z}_{j} through consideration:

$$\varphi_j = f(\operatorname{firm}_j, \operatorname{premium}_j, \operatorname{deductible}_j, \operatorname{gap}_j, \operatorname{Top100}_j, \operatorname{AvgCS}_j).$$
 (5.5)

To account for the higher weight placed on premiums relative to OOP drug costs in reduced form regressions of Section 4.3 and to capture potential behaviors stemming from unobserved constraints or presentation of plan information, a plan's premium may impact consideration. Premium enters consideration, however, in a different, non-linear way than utility (see below). The identity of the firm, the premium, the deductible, and whether or not a plan includes any supplemental gap coverage are all observable to a beneficiary searching for plans. To summarize the generosity of the plan formulary, I compute the count of the 100 most claimed drugs by spend among the Medicare population covered in each plan's formulary. As a simple measure of plan copays and coinsurance rates, I compute the average costshare in the initial coverage phase under each plan.²³ Although the illustrative example in Figure 3.1 was a simplification, the idea of such diminishing consideration is appealing and converges to a specification of consideration that reflects a geometric decay of consideration probabilities as plans progressively become less and less desirable in their attributes.

ASSUMPTION 5.2 (Consideration Specification):

1. As an intuitively appealing way to ensure consideration probabilities are in the unit interval, I impose the following functional form:

$$\varphi_j = \phi_{firm_j} \phi_{prem_j} \phi_{ded_j} \phi_{gap_j} \phi_{top100_j} \phi_{AvgCS_j}, \tag{5.6}$$

where $\phi_{firm_j} \in [0, 1]$ is the base consideration probability of the firm offering plan *j*, constant across all plans offered by that firm in the California market. The plan attributes enter consideration multiplicatively as well, with all δ terms $\in [0, 1]$

$$\begin{split} \phi_{prem_j} &= \delta_{prem}^{PremRatio}, \\ \phi_{ded} &= \delta_{ded}^{DedRatio} \\ \phi_{Gap} &= \begin{cases} \delta_{gap} & \text{if No Gap} \\ 1 & \text{if Gap}, \end{cases} \\ \phi_{Top100} &= \delta_{top100}^{(max(top100) - top100_j)}, \\ \phi_{AvgCS} &= \delta_{avgcS}^{(AvgCS_j - min(AvgCS))} \end{split}$$

Both ϕ_{prem_j} and ϕ_{ded_j} , which govern the roles of premium and deductible, respectively, depend on the ratio of a plan's premium and deductible relative to the max-

²³After estimating expenditures in every plan under perfect foresight, I compute the average ratio of out-of-pocket costs to gross expenditure within the initial coverage phase across all individual's whose out-of-pocket spending was between the deductible and initial coverage limit. Accordingly, this variable takes the same value for all beneficiaries.

imum in the market. I define $PremRatio_j \equiv \frac{Prem_j - min(Prem)}{max(Prem) - min(Prem)}$ and $DedRatio_j \equiv \frac{Deduc_j}{max(Deduc)}$. The consideration parameters to estimate include the 19 φ_{firm_j} terms, and the attribute δ terms.

5.3 Maximum Likelihood

The probability any beneficiary *i* selects plan j^* can be written as:

$$Pr(y_{ij^*} = 1) = \sum_{M \subseteq \mathcal{M}: j^* \in M} \prod_{k \in M} \varphi_k(\mathbf{Z}_k) \prod_{k' \notin M} \left(1 - \varphi_{k'}(\mathbf{Z}_{k'})\right) Pr\left(EU_{ij^*} \ge EU_{ik} \ \forall \ k \in M\right)$$

$$(5.7)$$

As written, equation 5.7 requires enumeration of all possible consideration sets M. In a setting such as Medicare Part D where beneficiaries in California have 46 plans available, such an enumeration is computationally infeasible. Rather than approximate such a sum with simulation of consideration sets, as done in Goeree (2008),²⁴ this choice probability can be simplified to fully avoid the need to account for every potential consideration set. The utility model in equation 5.4 does not include an error term, and at any given value of risk aversion ν_i , all plans can be ranked by expected utility. That is, fix $\hat{\nu}$, and order plans from worst to best in terms of expected utility $EU_{i1} < EU_{i2} < ... < EU_{ij^*} < EU_{ij+1}... < EU_{iJ}$. Therefore, for plan j^* to have been selected at $\nu = \hat{\nu}$, the consideration set must not have included (at the minimum) plans j +1, ..., J, since if those plans were present, j^* would not be selected. Let $k \succ_{\hat{\nu}} j^*$ denote the set of plans that dominate j^* at a given value $\hat{\nu}$. Thus, conditional on $\nu_i = \hat{\nu}$, $Pr(EU_{ij^*} > EU_{ik} \forall k \in M) = 0$ if M contains any plans in the set $k \succ_{\hat{\nu}} j^*$ and $Pr(EU_{ij^*} > EU_{ik} \forall k \in M) = 1$ if M does not contain any plans $k \succ_{\hat{\nu}} j^*$.

Such a ranking and collection of dominating plans can be computed at any value of $\nu \in [0, .\bar{\nu}]$, where $\bar{\nu}$ is the upper bound on the coefficient of absolute risk aversion. Using this simplification, equation 5.7 for a given value of ν_i can be written without regard for specific consideration set as:

²⁴Other studies focus only on the largest plans for computational tractability.

$$Pr(y_{ij^*} = 1 | \nu_i = \hat{\nu}) = \varphi_{j^*}(\mathbf{Z}_{\mathbf{j}^*}) \prod_{k \succ_{\hat{\nu}} j^*} (1 - \varphi_k(\mathbf{Z}_{\mathbf{k}}))$$
(5.8)

These sets of dominating plans can be computed for each individual at any value of risk aversion. Averaging equation 5.8 across individuals allows for approximation of the choice probabilities of the form:

$$Pr(y_{j^*} = 1) = \int Pr(y_{j^*}|\nu)dF_{\nu}.$$
(5.9)

In practice, the integral in equation 5.9 is estimated through a Riemann approximation. The support of the coefficient of risk aversion, $[0, \bar{\nu}]$, is divided into a fine grid. At each value of ν on the grid, for each individual, the set of plans $k \succ_{\nu} j^*$ is computed, as described in equation 5.8. To approximate the integral over the distribution of ν , I weight the choice probabilities above at each value of ν in the grid based on the probability density function of risk aversion at those grid values. Weighted individual choice probabilities are then logged and summed. I maximize the resulting loglikelihood to recover the values of the model parameters - including those governing the distribution of risk aversion - that best match the observed choices.

In all specifications, I assume the coefficient of absolute risk aversion follows a Beta distribution, $\nu \sim Beta(\beta_1, \beta_2)$. The Beta distribution is an appealing assumption due to its flexibility. Risk aversion is assumed to be bounded above by .01, a liberal assumption in light of Rabin (2000). Estimates are not sensitive to this assumption.

5.4 Identification

To separately identify consideration from risk preferences, I assume a large support of certain variables and a form of an exclusion restriction. There must be sufficient variation in the utility-relevant variables to shift utility rankings of plans without correspondingly shifting consideration probabilities. In this model, the only utility-relevant variables are those governing the distribution of costs under each plan. Other plan attributes are presumed to impact consideration but not directly enter utility, with the exception of premium which enters consideration nonlinearly. The considerationrelevant variables that relate to potential drug costs - the deductible, gap coverage, count of drugs covered, and realized average cost-share - are linked with the distribution of costs but in very complex, individual-specific manners. As a result, there is sufficient independence in variation between the plan attributes and the utility-relevant variables to argue exclusion. In fact a simple linear regression indicates that the consideration variables explain less than than 1.5% of the variation in the utility variables, with most of the variation coming from variation in individual drug needs.

Identification can be viewed in two stages. First, to identify the consideration probabilities, φ_j , I require a large support for the utility-relevant variables, $\hat{\mu}_{ij}$ and σ_{ij}^2 . Intuitively, there are regions of the support of these variables where the utility ranking of the plans does not depend on risk preferences. For example, regardless of risk aversion, in certain regions of the support a specific plan is unambiguously best, and under full consideration, I would expect to see all individuals in that region of the support choosing said best plan. The discrepancy between full enrollment and the empirical share of individuals selecting this plan identifies the consideration probability for that specific plan. In these regions of the support of the utility-relevant variables, choice probabilities can be written in terms of consideration probabilities alone. Provided there is an additional region of the support where plans again are ordered irrespective of preferences, but the order has now changed, there are sufficient moments to identify all plan consideration probabilities. Variation of plan attributes within and across firms identifies the individual components of consideration. The second step is to identify the distribution of risk preferences. With consideration identified, this proceeds in the same manner as a full consideration model. Variation in the mean and variance of costs traces out the distribution of ν among the population.

6 Results

The model of expected utility with limited consideration matches the data patterns of the California beneficiaries well. Heterogeneity in consideration sets plays a crucial role in explaining prescription drug insurance choice. Recall that individual choices arise from the interplay of two model primitives: the consideration set an individual draws from the available menu of plans and their risk preferences and the corresponding expected utility of each considered plan. Estimation yields two sets of results those governing consideration probabilities and those defining the distribution of risk preferences.

6.1 Limited Consideration

Each individual draws an unobserved consideration set M_i based on each plan's independent probability of consideration, φ_j . To capture the manner in which plan characteristics may impact the likelihood beneficiaries consider a given plan (for example if individuals are filtering the choice set by deductible), φ_j is the product of firm and attribute consideration probabilities. I first present the results on firm consideration probabilities and then those relating to plan attributes.

Figure 6.1 presents the insuring firm base consideration probabilities. The simplest way to interpret these values is to imagine a hypothetical plan that had the best of all attributes, meaning the lowest premium, zero deductible, gap coverage, highest count of covered drugs, and lowest out-of-pocket cost-sharing during the initial coverage phase. For such plan, the values in Figure 6.1 represent the plan consideration probability. Such a plan offered by Firm 1 would be in all consideration sets, and a plan offered by Firm 9 would be in around 20% of consideration sets. Of course such a plan does not exist, but these firm consideration probabilities represent the initial departure from full consideration that arises due to the impact of the identity of the insuring firm. Three large firms, UnitedHealth, Blue Cross of California, and Anthem, garner near full base consideration. Each of these firms offers a plan included in the 5 most chosen plans within the sample, and UnitedHealth offers an AARP-sponsored plan with the highest enrollment in the market. Nearly half of the 19 firms in the market are considered with probability below 10%, even before accounting for the impact of plan attributes. Such heterogeneous consideration across firms may reflect, among other explanations, the impact of differential advertising, agent steering effects, or enrollees' insurance experiences prior to Medicare. I leave to future research the detailed analyses of such explanations.

As noted, each plan also contains a collection of attributes that may impact the probability an individual considers it. Recall the illustrative example in Figure 3.1. As both deductible and premiums increase relative to the minimums available, consideration of those plans became less likely. As noted in Assumption 5.2 ensure consideration probabilities are in the unit interval, I impose:

$$\varphi_j = \phi_{firm_j} \phi_{prem_j} \phi_{ded_j} \phi_{gap_j} \phi_{top100_j} \phi_{AvgCS_j},$$



Figure 6.1: Model Results: Firm Base Consideration Probabilities

Notes: Firms are ordered based on estimated base consideration probabilities. Error bars present 95% confidence intervals based on 1,000 bootstrap repetitions with sub-sampling.

where $\phi_{firm_j} \in [0, 1]$ is constant across all plans offered by that firm in the California market. The plan attributes enter consideration multiplicatively as follows, with all δ terms $\in [0, 1]$

$$\begin{split} \phi_{prem_{j}} &= \delta_{prem}^{PremRatio}, \\ \phi_{ded} &= \delta_{ded}^{DedRatio} \\ \phi_{Gap} &= \begin{cases} \delta_{gap} \text{ if No Gap} \\ 1 \text{ if Gap}, \\ \phi_{Top100} &= \delta_{top100}^{(max(top100) - top100_{j})}, \\ \phi_{AvgCS} &= \delta_{avgcs}^{(AvgCS_{j} - min(AvgCS))} \end{split}$$

Table 6.1 presents the estimates of the impact of plan attributes on consideration. As described above, a plan's consideration probability is modeled as the product of a firm-specific base consideration probability and attribute impacts. The δ estimates reflect the *total* decay in consideration that occurs as the attribute changes from the most desirable value to the least desirable value. All else equal, the estimate for δ_{prem} in Table 6.1 indicates a plan with the highest premium is considered only 11.2% as often as the lowest premium plan. Similarly, a plan with the maximum deductible of \$310 is considered 17.4% as often as an equivalent zero deductible plan. For any value of deductible (premium) between the lowest and the highest, the δ term would be exponentiated based on the ratio of that plan's deductible (premium), as described above in Assumption 5.2. Plans lacking gap coverage are considered 84.7% as frequently as plans with gap coverage.

	Estimate	95% CI
δ_{prem}	0.112	[0.081, 0.156]
δ_{ded}	0.174	[0.154, 0.197]
δ_{gap}	0.847	[0.774, 0.929]
δ_{top100}	1.000	[1.000, 1.000]
δ_{avacs}	1.000	[1.000, 1.000]

Table 6.1: Model Results: Consideration Impact of Plan Attributes

Notes: All δ terms are defined between 0 and 1 and reflect how much consideration a plan with the worst value of an attribute is considered relative to an equivalent plan with the best value of the attribute. Confidence intervals based on 1,000 bootstraps with sub-sampling.

The results in Table 6.1 are rather intuitive. As modeled, a plan's count of top 100 drugs covered in the formulary and the average cost-share in the initial coverage phase do not impact its probability of consideration. These plan attributes are, generally, not immediately observable to a beneficiary. An individual can find whether certain drugs are covered in a plan's formulary through tools such as Medicare's PlanFinder online tool, but a full count of such coverage of the 100 most popular drugs among beneficiaries is not published. Additionally, an astute beneficiary that seeks out information on their plans can learn the copay and coinsurance rates for different tiers of drugs in the initial coverage phase - that information is precisely what the average cost-share variable is meant to proxy for - but such a precise aggregate measure is not easy to compute for most individuals. To the extent this captures filtering of choice sets based on desirable attributes, it is not particularly surprising that these more difficult to ascertain attributes are not strong drivers of consideration. The first three attributes, in contrast, are easily known to beneficiaries. In fact, many online tools, including the PlanFinder, summarize exactly this information for beneficiaries; individuals can even

sort and filter available plans on the PlanFinder by premiums and deductibles.²⁵

My estimates imply substantial heterogeneity in consideration sets. Figure 6.2 presents the implied distribution of consideration set sizes across individuals in the sample. Although the market includes 46 plans, consideration sets do not come anywhere close to including that many plans. The vast majority of beneficiaries consider no more than 5 plans, and no one is estimated to consider a set containing more than 14 plans. Approximately 15% of individuals evaluate a single plan. As shown in Figure 6.3, the largest firms account for an overwhelming share of the plans considered. Three large firms constitute over 60% of all plans considered. The three firms with the smallest firm base consideration probabilities account for below 0.5% of plans considered. In fact, 9 of the 19 firms each represent fewer than 1% of considered plans and cumulatively represent close to 4% of all plans considered. Figure 6.3 also emphasizes that this pattern is not an artifact of the number of plans offered, but rather, reflects the strong positioning of a few large firms.



Figure 6.2: Implied Distribution of Consideration Set Size

Notes: Consideration set sizes estimated as the average over 1,000 simulations of individual risk aversion and consideration sets using the analysis sample.

²⁵Following the Affordable Care Act, Medicare plans no longer include the coverage gap. As such, today's PlanFinder does not present this information.



Figure 6.3: Implied Shares of Consideration Sets and Choice Set by Firm

Notes: Firms are ordered as in Figure 6.1 based on estimated firm base consideration probabilities. Shares of consideration sets are based on 1,000 simulations of individual consideration sets using the analysis sample.

Consideration sets are similarly skewed towards zero deductible plans. Plans without a deductible account for 19 of the 46 plans offered in California in 2010 but nearly 67% of considered plans. Figure 6.4 illustrates this pattern in the first panel. The second panel of Figure 6.4 plots the share of considered plans based on bins of premiums. The pattern of the premiums of considered plans is more nuanced. The first bar represents the 10 lowest premiums, the second bar the next 10 lowest premiums, and so on. While the estimate of δ_{prem} conforms with the intuition that higher premium plans are considered less often than more appealing lower premium plans, the plans with the lowest premiums are generally those with higher deductibles. Thus, this preference towards lower premium plans alongside low deductibles manifests in the plans in the second bin of premiums accounting for a disproportionate share of plans considered.

The resulting consideration probabilities of the 46 available plans vary substantially. These consideration results are consistent with a number of underlying sources of limited consideration. The strong impact of the deductible on consideration coheres with stories of liquidity constraints, a reality for many Americans, as noted in Durante and Chen (2019). It is both plausible and rational for such a constrained beneficiary



Figure 6.4: Implied Shares of Consideration Sets and Choice Set by Deductible and Premium

(a) Share of Consideration Sets by Deductible (b) Share of Consideration Sets by Premium (bin)

Notes: Panel (a) is ordered left to right from \$0 to \$310 deductibles. Premiums in Panel (b) are ordered lowest to highest by bins of 10. Shares of consideration sets are based on 1,000 simulations of individual consideration sets using the analysis sample.

to consider exclusively, or nearly exclusively, plans with an eliminated deductible, as my estimates indicate. The substantial role of firm effects in consideration lends support to a number of behavioral forces resulting in limited consideration. Familiarity of firms based on prior insurance experience or social influence, such as the insurer of the beneficiary's friends or spouses, can lead beneficiaries to filter available plans according to preferred firms. Alternatively these results may reflect the consideration impact of advertising campaigns.

The overall impact of these mechanisms results in consideration sets that are much smaller in size than the feasible choice set. The modal consideration set contains 3 plans and nearly 94% of beneficiaries consider a set with 5 or fewer plans. To parse the effects of each attribute on the resulting consideration set composition, I simulate consideration sets when certain sources of consideration are eliminated. Figure 6.5 plots the distribution of consideration set sizes across two such schemes. Holding all other estimates fixed, Panel (a) presents the impact on consideration set size of eliminating the firm effect. In practice, this exercise translates to assigning all firms a base consideration probability of 1. This alteration results in a rightward shift of the distribution of consideration set size, as fewer plans are immediately eliminated as a result of firm filtering. Panel (b) presents the opposite exercise where the firm effect is the sole determinant of consideration. Because the three largest firms have base consideration probabilities of 1, or nearly 1, by construction consideration sets have a larger minimum number of plans. Similar, but less drastic patterns emerge when premium, deductible, and gap coverage are shut down.

Figure 6.5: Baseline and Counterfactual Distributions of Consideration Set Size



Notes: Implied consideration sets sizes of the baseline in blue and the following adjustments simulated in red: Panel (a) all firm base probabilities are set to 1; Panel (b) δ_{prem} , δ_{ded} , and δ_{gap} in Table 6.1 are all set to 1; Panel (c) δ_{prem} is set to 1; Panel (d) δ_{ded} is set to 1. Shares of consideration sets are based on 1,000 simulations of individual consideration sets.

6.2 Risk Preferences

Upon drawing a consideration set, beneficiaries evaluate plans based on the expected utility implied by their level of risk aversion. In contrast to the previous literature on plan choice in Medicare Part D, I find estimates of moderate risk aversion among California's beneficiaries comparable to other insurance settings. In the Part D existing literature, risk preferences were either not the focus of estimation or found to be largely irrelevant.²⁶ Risk aversion is modeled to be heterogeneous across beneficiaries and follows a Beta distribution. Table 6.2 presents estimates of risk preferences in the model with limited consideration, the result of Beta parameter estimates of (0.7, 11.3). The estimate of mean risk aversion is on par with, although lower than, previous studies that

 $^{^{26}}$ Note, as many studies estimate a linear approximation of expected utility, it is not immediately clear how much the difference might be due to limited consideration versus the effects of linear approximation.

use field data to measure risk preferences, and comes alongside moderate variance.²⁷

These estimates can be difficult to interpret and compare without additional context. Table 6.2 includes a measure of risk premium for an individual with CARA utility facing a lottery that results in a loss of \$1,000 with 25% probability. Such a lottery has an expected value of a \$250 loss. An individual with a coefficient of risk aversion equal to my baseline mean estimate would be willing to pay a risk premium of \$57 to avoid such a lottery. In contrast, a standard CARA random expected utility model with full consideration substantially underestimates risk aversion, as shown in the second column of Table 6.2, with the upper bound of the confidence interval below the lower bound of the Baseline model for both the mean and variance of risk aversion. A CARA full consideration model with a homogeneous coefficient of risk aversion across agents finds risk neutrality and matches choice patterns very poorly. The inclusion of non-monetary and non-cost attributes is important for estimating risk aversion in this setting.

Table 6.2: Model Estimates: Risk Preferences

	Baseline	CARA RUM (heterogeneous)
$\mathbb{E}(\text{Risk Aversion})$	$5.58 \cdot 10^{-4}$	$6.18\cdot10^{-5}$
	$[3.51 \cdot 10^{-4}, 9.18 \cdot 10^{-4}]$	$[5.73 \cdot 10^{-5}, 6.73 \cdot 10^{-5}]$
Var(Risk Aversion)	$4.07 \cdot 10^{-7}$	$1.85 \cdot 10^{-8}$
	$[1.02 \cdot 10^{-7}, 2.01 \cdot 10^{-6}]$	$[1.72 \cdot 10^{-8}, 2.01 \cdot 10^{-8}]$
Risk Premium	\$57	\$6

Notes: Based on 1,000 bootstraps with sub-sampling in Baseline model. CARA RUM assumes full consideration. Risk premium is calculated at mean risk aversion for a lottery with a 25% probability of a loss of \$1,000.

6.3 Plan Choice under Limited Consideration

Choices in this market are driven both by risk aversion and limited consideration. Taken together, the model of expected utility with limited consideration matches the observed choice patterns of beneficiaries well. Figure 6.6 plots the implied choice distribution of the baseline model alongside the empirical distribution of plan choices.

²⁷See Barseghyan et al. (2021a), Handel and Kolstad (2015), Handel (2013), Barseghyan et al. (2013), and Cohen and Einav (2007). In particular, Barseghyan et al. (2021a) finds that with unobserved choice sets can rationalize auto collision choices with lower and more homogeneous risk aversion than standard models.



Figure 6.6: Empirical and Model Choice Distributions

Notes: The blue bars correspond to the empirical choice shares, and the red bars are the implied choice distribution based on 1,000 simulations. Plans are ordered from the plan with the largest enrollment share on the left to the plans with zero enrollment. The rightmost plan corresponds to a composite plan of the 17 plans in which between 1 and 10 individuals enrolled.

I am able to simulate choices under full consideration using the estimated distribution of risk aversion. As a result of limited consideration, individuals cluster on plans with low deductibles offered by a few popular firms. Many of these plans, however, are not particularly well suited to the drug needs and risk preferences of many individuals. As a measure of the cost of limited consideration in this population, I compute the difference in certainty equivalent of the chosen plan and the preference-implied optimal plan. The average certainty equivalent loss across the sample is \$223. Table 6.3 compares the distribution of certainty equivalent difference under the same exercise in Figure 6.5 in which consideration effects are eliminated. The reduction in the average difference in certainty equivalent between chosen and optimal plans arises from the corresponding increase in consideration set sizes displayed in Figure 6.5. The changes in certainty equivalent difference highlight the sizable but varying role of attributes in consideration and the cost of limited consideration. Additionally, as seen in the estimates, there is sizable heterogeneity across individuals in the cost of limited consideration.

Since the model of limited consideration nests many possible underlying sources of consideration set formation, the model is agnostic about why the attributes determine consideration sets. Take, for example, the firm effect. If the impact of firm identity on choices can be purely a consideration impact; advertising by some firms in this market is substantial. In that case, the difference between the baseline certainty equivalent loss of \$223 and the \$85 loss under a scenario without the firm effect represents meaningful welfare improvement. By limiting the set of plans considered based on firm identity, beneficiaries are on average leaving \$138 on the table. If, however, if the impact of firm on plan consideration represents perceived quality, then a situation that removes an individual's ability to filter plans considered based on the firm would be welfare reducing. In that case the \$138 additional lost certainty equivalent can be interpreted as an average bound on the shadow price of unobserved quality.

A similar logic applies to the other results. If the role of the deductible is through consideration, then removing this channel and the corresponding increase in consideration set sizes improves welfare.²⁸ Individuals save on average \$87 in that counterfactual. However, if the role of the deductible in consideration represents liquidity constraints, then individuals will become worse off without the ability to limit plans based on deductible. For the deductible, premium, and gap effects, if the impact my model captures is consideration rather than constraints, the changes in average certainty equivalent can be interpreted as clear welfare improvements through reducing consideration obstacles. If, however, these effects are manifestations of binding constraints, these changes represent the shadow prices of said constraints.

6.4 Counterfactual Exercises

There has been discussion since Part D's inception that the large number of available plans is unmanageable even for the most sophisticated of enrollees. Changes to market regulations over time have reduced the number of available plans from approximately 50 in the early years of the program to closer to 30 plans today. My model does not impose a specific mechanism behind limited consideration. As such, there are meaningful limitations on credible counterfactual analyses under such agnosticism. To estimate the impact of such a large, or even larger, reductions in the size of the feasible set makes strong assumptions about the portability of my consideration estimates. It seems plausible that my reduced form estimates of consideration, presumably averaging

²⁸For example, available plans in the CMS sponsored online PlanFinder tool can be sorted by deductible. It is plausible an individual presented with plans in that order does not look beyond the initially presented plans and only considers zero or low deductible plans.

	Percentile					
	Mean	5^{th}	25^{th}	50^{th}	75^{th}	95^{th}
Baseline Model	223	43	125	187	284	499
Removing Firm Effect	85	14	41	70	104	210
Only Firm Effect	84	5	35	65	109	224
Removing Deductible Effect	136	23	72	103	171	331
Removing Premium Effect	153	20	71	135	194	367
Removing Gap Effect	196	35	106	162	250	448

Table 6.3: Baseline and Eliminating Consideration Results: Certainty Equivalent (\$) Loss Due to Limited Consideration

Notes: Average CE difference is the difference in certainty equivalent of the chosen plan and the optimal plan over 1,000 simulations of risk aversion and consideration sets. The counterfactual values are computed similarly when consideration sets are simulated with different components of plan consideration probabilities set to 1.

over different types of agent behavior, would differ in a substantially altered choice environment where individuals may not need to rely on attributes so heavily to navigate the choice environment.

As such, I estimate four counterfactuals varying either the feasible menu or the consideration process. I first conduct a counterfactual in which plans that are of average lower utility are dropped from the menu, as in Brown and Jeon (2020). In this case, I drop the 11 plans in the lowest quartile of average utility and simulate consideration sets and plan choice under this reduced menu. The second exercise simulates a marginal change to consideration sets, assessing the potential certainty equivalent gain from ensuring that each individual considers at least one more plan as a measure of the marginal impact of consideration. I do so by adding a random plan from the subset of 23 plans (top half) that on average deliver the highest utility. This highlights the value of ensuring that each individual considers at least one good plan, where good indicates a plan that is on average in the top half of utility. Finally, as a reference point for all results, I simulate choices from assigning each person randomly to a single plan, thus eliminating the role of preferences entirely from choice. The random assignment is performed both from the entire feasible set as well as from the half of plans with the highest average utility.

Table 6.4 presents the distributions of certainty equivalent difference between chosen and optimal plans. A few takeaways are immediately apparent. Individuals fare much worse under random assignment than in the baseline, although randomly assigning individuals to a large group of relatively good plans lessens the potential welfare loss as expected. The change in potential welfare of adding just one additional plan to an individual's consideration set is substantial. The average individual in that scenario recovers \$54 in CE loss, around a quarter of the baseline estimate, and this change is relatively constant across the entire distribution of CE losses. This highlights the substantial potential welfare improvements from even marginal changes in consideration sets.

Certainty Equivalent (\$) Changes Percentile 5^{th} 25^{th} 50^{th} 95^{th} 75^{th} Mean Baseline Model 22343125284499 187Dropping Lowest Quartile of Plans 216 41 119177280498One Additional Plan - Top Half of Plans 32 395 16991136216300 904 44895 403558Random Assignment Random Assignment - Top Half of Plans 338 53187 259441 812

Table 6.4: Counterfactual Exercises:

Notes: Average CE Difference computed as the difference in certainty equivalent of the chosen plan and the optimal plan over 1,000 simulations of risk aversion and consideration sets. Note in the case of dropping the lowest quartile of plans, CE is measured relative to the optimal plan in the remaining available plans.

Conclusion 7

Since prescription drug coverage was introduced to Medicare in 2006, researchers have encountered challenges in rationalizing a sizable fraction of observed plan choices, as is the case in many other health insurance markets. Expected utility alone, the classic workhorse model of insurance choices, does not match the choices of beneficiaries well. Alternative methods of adding non-cost plan attributes, which are important for matching empirical patterns, into a utility framework result in estimates that are difficult to structurally interpret in this environment. There are numerous plausible explanations for limited consideration in this market. Even for individuals lacking cognitive limitations, the time required to consider and compare so many plans may simply be too costly. These individuals may use certain plan attributes to trim the

choice set down to a manageable size. The reduction of the choice set according to attributes may reflect unobserved constraints on an individual or the impact of firm advertising or the presentation of plans to the beneficiary. My model does not specify the underlying source of limited consideration but provides important insight into what features of plans drive consideration. Moreover, my model of consideration and the relationship of plan attributes to consideration, is computationally tractable and is not subject to a curse of dimensionality as feasible choice sets increase in size. This feature is especially appealing as the vector of estimated parameters may converge to a fixed number as the size of the choice set is increased. In a market with many choice sets, including the PDP market, this is a very useful model feature.

My results show that heterogeneity in consideration sets plays an important role in plan choice. Despite the set of available plans, beneficiaries are largely considering no more than 5 plans. Results additionally highlight the importance of accounting for consideration when estimating risk preferences. I find estimates of mean risk aversion more than twice that of a full consideration model. My results on preferences and consideration are informative to policymakers as Americans increasingly encounter choices over health insurance plans. With influence over product standardization, presentation of information, and firm behavior, policymakers may be able to harness the information about how beneficiaries are choosing prescription drug plans to help remove the obstacles that prevent so many beneficiaries from considering and choosing their optimal plan in terms of utility. Overall, my results suggest that the existing evidence documenting the inconsistencies between Medicare Part D plan choice and standard economic models of insurance demand may not be as conclusive as previously thought. By incorporating limited consideration, individual choice behavior comports with an otherwise rational expected utility model of insurance choice. The results also suggest, however, that care should be taken when designing the environment in which consumers select their own insurance plans.

Acknowledgments: I am grateful to my dissertation committee, Levon Barseghyan, Francesca Molinari, and Colleen Carey for their guidance and encouragement. This paper has also benefited from conversations with Ben Handel, Matthew Thirkettle, Panle Jia Barwick, Giulia Brancaccio, Lin Xu, and participants in various conferences and seminars. This research was supported, in part, by the L.R. "Red" Wilson endowment.

References

- ABALUCK, J. AND A. ADAMS-PRASSL (2021): "What do consumers consider before they choose? Identification from asymmetric demand responses," *The Quarterly Journal of Economics*, 136, 1611–1663.
- ABALUCK, J. AND J. GRUBER (2011): "Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D program," *American Economic Review*, 101, 1180–1210.
- (2016a): "Choice inconsistencies among the elderly: Evidence from plan choice in the Medicare Part D program: Reply," *American Economic Review*, 106, 3962–87.
 (2016b): "Evolving choice inconsistencies in choice of prescription drug insurance," *American Economic Review*, 106, 2145–84.
- APESTEGUIA, J. AND M. A. BALLESTER (2018): "Monotone stochastic choice models: The case of risk and time preferences," *Journal of Political Economy*, 126, 74–106.
- BARSEGHYAN, L., M. COUGHLIN, F. MOLINARI, AND J. C. TEITELBAUM (2021a): "Heterogeneous choice sets and preferences," *Econometrica*, 89, 2015–2048.
- BARSEGHYAN, L., F. MOLINARI, T. O'DONOGHUE, AND J. C. TEITELBAUM (2013): "The Nature of Risk Preferences: Evidence from Insurance Choices," *American Economic Review*, 103, 2499–2529.
- BARSEGHYAN, L., F. MOLINARI, AND M. THIRKETTLE (2021b): "Discrete choice under risk with limited consideration," *American Economic Review*, 111, 1972–2006.
- BHARGAVA, S., G. LOEWENSTEIN, AND J. SYDNOR (2017): "Choose to lose: Health plan choices from a menu with dominated option," *The Quarterly Journal of Economics*, 132, 1319–1372.
- BROWN, Z. Y. AND J. JEON (2020): "Endogenous information and simplifying insurance choice," Tech. rep., Mimeo, University of Michigan.
- COHEN, A. AND L. EINAV (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97, 745–788.
- DIEBOLD, J. (2016): "The effects of medicare part D on health outcomes of newly covered medicare beneficiaries," *The Journals of Gerontology: Series B*, 73, 890– 900.
- DURANTE, A. AND L. CHEN (2019): "Report on the economic well-being of US households in 2018 [Report for the Board of Governors of the Federal Reserve System]. Washington, DC: Board of Governors of the Federal Reserve System,".
- ERICSON, K. M. AND A. STARC (2012): "Heuristics and heterogeneity in health insurance exchanges: Evidence from the Massachusetts connector," *American Economic*

Review, 102, 493–97.

- ERICSON, K. M. AND J. R. SYDNOR (2018): "Liquidity constraints and the value of insurance," Tech. rep., National Bureau of Economic Research.
- ERICSON, K. M. M. AND A. STARC (2016): "How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics*, 50, 71–85.
- FOSGERAU, M., E. MELO, A. D. PALMA, AND M. SHUM (2017): "Discrete choice and rational inattention: A general equivalence result," *Available at SSRN 2889048*.
- GOEREE, M. S. (2008): "Limited information and advertising in the US personal computer industry," *Econometrica*, 76, 1017–1074.
- HANDEL, B. R. (2013): "Adverse selection and inertia in health insurance markets: When nudging hurts," *American Economic Review*, 103, 2643–82.
- HANDEL, B. R. AND J. T. KOLSTAD (2015): "Health insurance for" humans": Information frictions, plan choice, and consumer welfare," *American Economic Review*, 105, 2449–2500.
- HEISS, F., A. LEIVE, D. MCFADDEN, AND J. WINTER (2013): "Plan selection in Medicare Part D: Evidence from administrative data," *Journal of Health Economics*, 32, 1325–1344.
- HEISS, F., D. MCFADDEN, AND J. WINTER (2010): "Mind the gap! Consumer perceptions and choices of Medicare Part D prescription drug plans," in *Research* findings in the economics of aging, University of Chicago Press, 413–481.
- HO, K., J. HOGAN, AND F. SCOTT MORTON (2017): "The impact of consumer inattention on insurer pricing in the Medicare Part D program," *The RAND Journal of Economics*, 48, 877–905.
- HORTAÇSU, A., S. A. MADANIZADEH, AND S. L. PULLER (2017): "Power to choose? An analysis of consumer inertia in the residential electricity market," *American Economic Journal: Economic Policy*, 9, 192–226.
- KAISER FAMILY FOUNDATION THE HARVARD School OF AND Public and theHEALTH (2006): Seniors Medicare Prescrip-Drug Benefit, https://www.kff.org/medicare/poll-finding/ tion seniors-and-the-medicare-prescription-drug-benefit/.
- KEANE, M. P., J. D. KETCHAM, N. V. KUMINOFF, AND T. NEAL (2019): "Evaluating Consumers' Choices of Medicare Part D Plans: A Study in Behavioral Welfare Economics," Tech. rep., National Bureau of Economic Research.

KESTERNICH, I., F. HEISS, D. MCFADDEN, AND J. WINTER (2013): "Suit the

action to the word, the word to the action: Hypothetical choices and real decisions in Medicare Part D," *Journal of Health Economics*, 32, 1313–1324.

- KETCHAM, J. D., N. V. KUMINOFF, AND C. A. POWERS (2016): "Choice inconsistencies among the elderly: Evidence from plan choice in the Medicare Part D program: Comment," *American Economic Review*, 106, 3932–61.
- (2019): "Estimating the Heterogeneous Welfare Effects of Choice Architecture," International Economic Review.
- KETCHAM, J. D., C. LUCARELLI, E. J. MIRAVETE, AND M. C. ROEBUCK (2012): "Sinking, swimming, or learning to swim in Medicare Part D," *American Economic Review*, 102, 2639–73.
- KETCHAM, J. D., C. LUCARELLI, AND C. A. POWERS (2015): "Paying attention or paying too much in Medicare Part D," *American Economic Review*, 105, 204–33.
- KLING, J. R., S. MULLAINATHAN, E. SHAFIR, L. C. VERMEULEN, AND M. V. WROBEL (2012): "Comparison friction: Experimental evidence from Medicare drug plans," *The Quarterly Journal of Economics*, 127, 199–235.
- LUCARELLI, C., J. PRINCE, AND K. SIMON (2012): "The welfare impact of reducing choice in Medicare Part D: A comparison of two regulation strategies," *International Economic Review*, 53, 1155–1177.
- MANSKI, C. F. (1977): "The structure of random utility models," *Theory and decision*, 8, 229–254.
- MANZINI, P. AND M. MARIOTTI (2014): "Stochastic choice and consideration sets," *Econometrica*, 82, 1153–1176.
- MATĚJKA, F. AND A. MCKAY (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105, 272–98.
- POLYAKOVA, M. (2016): "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D," *American Economic Journal: Applied Economics*, 8, 165–95.
- RABIN, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281–1292.
- SEMILLA, A. P., F. CHEN, AND T. M. DALL (2015): "Reductions in mortality among Medicare beneficiaries following the implementation of Medicare Part D." *The American journal of managed care*, 21, s165–71.
- SIMON, H. A. (1959): "Theories of decision-making in economics and behavioral science," The American economic review, 49, 253–283.

Online Appendix

Appendix A Cost Calculator and Distribution Estimation

Using detailed data on the plan cost structure, I construct a program to compute the out-of-pocket costs for any series of ordered drug claims under every available plan in 2010 in the 34 major regions in the United States. The detailed claims data include information about gross and out-of-pocket realized drug costs under chosen plans, but a cost calculator is required to compute the counterfactual drug expenditures under the plans individuals did not select. The first step of the calculator is the collect the relevant set of plans to construct costs based on CMS region.

A.1 Drug Cost Calculator

For any sequence of claims, I identify the coverage classification of each drug under each available plan. There are multiple numeric codes used to identify drug by molecule, formulation, and strength across the relevant data sets. These numeric systems do not, however, identify drugs uniquely. The claims data identifies drugs by National Drug Codes (NDCs), as well as a CMS created number referred to as the Formulary RX ID. The public use formulary data identify drugs by NDCs and RXCUIs. However, multiple NDCs can be used for the same drug. As such, NDCs are considered the same if they are linked through Formulary RX IDs. For example, consider a drug denoted as NDC_1 and $FRXID_1$. If NDC_1 is also listed as corresponding to $FRXID_2$, and $FRXID_2$ is elsewhere linked to NDC_2 , I consider NDC_1 and NDC_2 the same drug. For each claim passed through the calculator, I apply the lowest tier and base price of any linked NDC, allowing for some potential, albeit minor, substitution.

For every drug included in a plan's formulary, I determine the tier of coverage and whether that tier is covered in the donut hole. I also determine the base price of the claim by scaling the negotiated price of a 30 days supply of each drug under each plan to correspond to the days supply claimed. My calculator then processes the claims sequentially, determining the coverage phase and applying the relevant cost-sharing based on tier and phase. In the event a claim straddles multiple coverage phases, I prorate the claim across spending zones in the manner CMS does in practice. The calculator keeps a running total of gross and out-of-pocket spending throughout the series of claims.

To assess the performance of the cost calculator, I compare the estimated outof-pocket spending for each beneficiary's chosen plan to their realized out-of-pocket costs in the claims data. When I use as the base price of a drug the gross cost listed in the claims data, predicted and observed out-of-pocket spending have a correlation in excess of .95 for individuals across regions. In practice, I use the negotiated base prices listed in the pricing data to account for differences in base prices across plans. Occasionally there is a discrepancy between the information in the pricing file and what is reported in the claims data. Once I incorporate the negotiated base prices, the correlation between predicted and observed spending is .93 among the analysis sample. This simple test of accuracy is assuring, especially as I made a number of small simplifications in constructing the calculator that would prevent perfect prediction. I treated all claims as filed through in-network pharmacies and pro-rated one month cost-sharing for tractability.¹ In the catastrophic coverage phase, I treat all claims as though they are branded drugs. Out-of-pocket costs for those claims are therefore computed as the maximum between a 5% coinsurance and a \$6.30 copay. In practice, for generic drugs, the beneficiary pays the maximum of a 5% coinsurance and a \$2.50 copay. The data I use does not include information on whether a drug is branded or generic. However, few individuals enter the catastrophic coverage phase at all, and the differences in cost between these two pricing schemes is small.

A.2 Distribution of Out-of-Pocket Costs and Variance Estimation

As described in Section 4.2.2, higher order moments - and the mean under an assumption of rational expectations - of the distribution of drug costs an individual expects under different plans requires an approximation of the distribution of out-of-pocket drug costs under each available plan. In practice, it is the variance of costs for which I need estimates. To this end, I assign each individual in my sample into a bin of "similar individuals" based on their average monthly gross drug costs and average monthly "effective" claim counts during their 2010 tenure. Effective claim counts adjusts counts for the number of months a claim covers. For example, if a beneficiary filled a claim

 $^{^1\}mathrm{Typically},$ the three month copay was simply 3 times the one month copay, making this simplification innocuous.

for a 90 day supply, it is treated as effectively 3 claims. Average claim counts are are classified as one of the following: between 0 and 1, between 1 and 2, between 2 and 3, between 3 and 4, between 4 and 10, and more than 10. These claims bins are crossed with quintiles of average monthly gross spend. An additional bin of individuals with zero claims and zero spend is also defined. Bins with fewer than 100 individuals are dropped.

To estimate the cost distribution within each bin, I construct a sample of individuals without ESRD who are enrolled in a Part D plan for some portion of 2010 and the entirety of 2011. I use their 2010 claims experience to categorize them into one of the bins described above. I then randomly select 100 individuals from each bin and pass their entire 2011 claims experiences through the cost calculator for every plan. I compile monthly running totals of out-of-pocket spend for each randomly sampled individual. To adjust for the evolution of drug expenditure over time, I deflate all 2011 costs by the average ratio of 2010 spend compared to 2011. For every individual in my analysis sample, denoting their months of 2010 coverage by m, the variance of out-of-pocket costs in each plan is computed as the variance of the random sample's deflated out-of-pocket costs for m months of 2011.

Appendix B Robustness Analysis

B.1 Alternative Assumptions on OOP Costs

In my baseline analysis I assume perfect foresight of the mean of out-of-pocket drug costs. As a robustness check, I estimate my model using alternative specifications of expected out-of-pocket costs. I consider three cases. In the first case, individuals have rational expectations over costs, and μ_{ij} is computed in the same binning procedure as σ_{ij}^2 . I then consider two cases in which individuals are myopic and assume their current prescription drug use will continue throughout their coverage year, where because I lack information on drug claims at the time of enrollment I project out the first month or first two months of claims for the entire coverage year when computing μ_{ij} .

Table B.1 presents estimates of the consideration impact of plan attributes. The overall results are consistent. In the case of rational expectations, all else equal, these estimates suggest the highest premium plan is considered only 9% as much as the lowest premium plan. Plans with \$310 deductibles receive approximately 20% as much

consideration as equivalent \$0 plans. The impact of gap coverage is slightly more pronounced. As in the baseline analysis, the count of top 100 drugs covered and the average cost-share of a plan do not impact consideration. The last two columns present the results under the myopic approaches. The results are very similar to the baseline model, with somewhat more stark results for deductible and premium. Figure B.1 plots firm base consideration probabilities in the same manner as Figure 6.1 for the baseline case and each of the three robustness analyses. The same patterns emerge as in the baseline results. Additionally estimates of the mean and variance of risk aversion are quite stable across specifications.

Estimates of risk aversion are similar in all three robustness analyses to the baseline analysis. Table B.2 presents estimates and confidence intervals for the mean and variance of risk aversion. There is substantial overlap in the confidence intervals for both statistics. These estimates show that the results in the baseline analysis are not driven by the assumption of perfect foresight of expected out-of-pocket expenses.

Baseline	Rational Expectations	Myopic - 1 month	Myopic - 2 months
0.112	0.092	0.078	0.089
[0.081, 0.156]	[0.066, 0.130]	[0.057, 0.115]	[0.063, 0.128]
0.174	0.195	0.113	0.130
[0.154, 0.197]	[0.173, 0.221]	[0.099, 0.129]	[0.115, 0.147]
0.847	0.744	0.896	0.872
[0.774, 0.929]	[0.678, 0.834]	[0.814, 0.986]	[0.796, 0.956]
1.000	1.000	1.000	1.000
[1.000, 1.000]	[1.000, 1.000]	[1.000, 1.000]	[1.000, 1.000]
1.000	1.000	1.000	1.000
[1.000, 1.000]	[1.000, 1.000]	[1.000, 1.000]	[1.000, 1.000]
	Baseline 0.112 [0.081, 0.156] 0.174 [0.154, 0.197] 0.847 [0.774, 0.929] 1.000 [1.000, 1.000] 1.000 [1.000, 1.000]	BaselineRational Expectations0.1120.092[0.081, 0.156][0.066, 0.130]0.1740.195[0.154, 0.197][0.173, 0.221]0.8470.744[0.774, 0.929][0.678, 0.834]1.0001.000[1.000, 1.000][1.000, 1.000]1.0001.000	BaselineRational ExpectationsMyopic - 1 month0.1120.0920.078[0.081, 0.156][0.066, 0.130][0.057, 0.115]0.1740.1950.113[0.154, 0.197][0.173, 0.221][0.099, 0.129]0.8470.7440.896[0.774, 0.929][0.678, 0.834][0.814, 0.986]1.0001.0001.000[1.000, 1.000][1.000, 1.000]1.000[1.000, 1.000]1.0001.000

Table B.1: Robustness Results: Consideration Impact of Plan Attributes

Notes: All δ terms are defined between 0 and 1 and reflect how much consideration a plan with the worst value of an attribute is considered relative to an equivalent plan with the best value of the attribute. Confidence intervals are shown in brackets and based on 1,000 bootstraps with sub-sampling.

It is noted in the body of the text that the assumption of how individuals form expectations for costs is far from settled in the literature and alternatives beyond those presented here are certainly plausible. Without additional data or different as-



Figure B.1: Robustness Results: Firm Base Consideration Probabilities

Notes: Firms are ordered based on estimated base consideration probabilities in the baseline model, as in Figure 6.1. Error bars present 95% confidence intervals based on 1,000 bootstrap repetitions with sub-sampling.

	Baseline	Rational Expectations
$\mathbb{E}(\text{Risk Aversion})$	$5.58 \cdot 10^{-4}$	$1.39 \cdot 10^{-4}$
	$[3.51 \cdot 10^{-4}, 9.18 \cdot 10^{-4}]$	$[1.77 \cdot 10^{-13}, 5.28 \cdot 10^{-4}]$
Var(Risk Aversion)	$4.07 \cdot 10^{-7}$	$2.50\cdot 10^{-7}$
	$[1.02 \cdot 10^{-7}, 2.01 \cdot 10^{-6}]$	$[2.66 \cdot 10^{-16}, 3.07 \cdot 10^{-6}]$
Risk Premium	\$57	\$13
	Myopic - 1 month	Myopic - 2 months
$\mathbb{E}(\text{Risk Aversion})$	$6.11\cdot 10^{-4}$	$6.09\cdot 10^{-4}$
	$[3.73 \cdot 10^{-4}, 8.77 \cdot 10^{-4}]$	$[3.83 \cdot 10^{-4}, 9.55 \cdot 10^{-4}]$
Var(Risk Aversion)	$3.46 \cdot 10^{-7}$	$3.12 \cdot 10^{-7}$
	$[1.07 \cdot 10^{-7}, 9.10 \cdot 10^{-7}]$	$[1.08 \cdot 10^{-7}, 1.49 \cdot 10^{-6}]$
Risk Premium	\$63	\$63

Table B.2: Robustness Estimates: Risk Preferences

Notes: CI based on 1,000 bootstraps with sub-sampling. Risk premium is calculated for a beneficiary facing a lottery that results in a loss of \$1,000 with 25% probability.

sumptions it is hard to determine which approach best approximates reality, but these additionally analyses are reassuring that the main results do not appear to be driven by the assumption of perfect foresight.

B.2 Full Consideration and Alternative Models

The workhorse model of insurance demand is the expected utility model with full consideration. Similar to the model under limited consideration presented in this paper, the only utility-relevant variables are those governing the distribution of losses. Nonmonetary attributes are not provided a role in the decision framework. Beginning with the expected utility specification in Equation 5.4, I estimate a random utility model for comparison and to highlight the empirical advantages of accounting for limited consideration.² In both the homogeneous and heterogeneous risk aversion specifications,

 $^{^{2}}$ It is worth noting that Apesteguia and Ballester (2018) describe a theoretical shortcoming of the random expected utility model in the insurance setting due to the implied non-monotonicity of choice probabilities in risk aversion. The limited consideration model I employ avoids this issue.

the utility error is assumed to be iid Type 1 Extreme Value distributed.

$$EU_{ij} = -\exp(\nu\hat{\mu}_{ij} + \frac{1}{2}\nu^2\sigma_{ij}^2) + \epsilon_{ij}$$
(B.1)

$$EU_{ij} = -\exp(\nu_i \hat{\mu}_{ij} + \frac{1}{2}\nu_i^2 \sigma_{ij}^2) + \epsilon_{ij}, \ \nu_i \sim Beta(\beta_1, \beta_2)$$
(B.2)

The resulting estimates of risk preferences are described in Table 6.2 in the main text. The assumption of full consideration in the CARA expected utility model results in substantial underestimation of risk aversion. In a model of homogeneous preferences, beneficiaries are estimated to be effectively risk neutral, with an estimated risk premium for a 25% loss of \$1,000 of approximately a nickel. This is a puzzling result in an insurance market, and the implied choice probabilities of this model come close to rolling a 46-sided die. The inclusion of random preferences also underestimates risk aversion relative to the model of limited consideration. The omission of non-monetary, and non-cost more generally, attributes from a model of plan choice also diminishes the ability of the model to rationalize observed choice patterns, with this model seriously underestimating the most popular plan and greatly overstimating the share of plans with very low enrollment.

Acknowledging the importance of non-monetary attributes in rationalizing the choice of prescription drug plans, previously used methods take the approach of adding the plan attributes directly into utility. This can be done by scaling up the attributes by a coefficient to estimate. This suggests an interpretation of the coefficient as translating the variable into a "cost", comparable to the monetary attributes such as premium and out-of-pocket costs. Denoting the included non-monetary attribute by X_j and the monetizing scaling coefficient as γ , this translates in its simplest form here to the modified expected utility specification:

$$EU_{ij} = -\exp(\nu_i(\hat{\mu}_{ij} + \gamma X_j) + \frac{1}{2}\nu_i^2 \sigma_{ij}^2)$$
(B.3)

In practice, expected utility is often estimated as a conditional logit by including the non-monetary attributes additively and a Type 1 Extreme Value error.³ The estimates of coefficients on non-monetary attributes are generally interpreted relative

³See Abaluck and Gruber (2011) for a derivation of the conditional logit as a linear approximation of a CARA expected utility model.

to the coefficient on either premium or out-of-pocket costs as a willingness to pay for the attribute. Such a comparative interpretation is common in discrete choice models. As discussed in Handel and Kolstad (2015) and Handel (2013), this approach treats the utility cost of the attribute as constant across the distribution of losses. Effectively, the utility cost of the attribute is a mean shift of the distribution of drug costs arising from uncertain drug needs. It can be difficult to attribute an economic meaning to these estimates in some settings. If the inclusion of a non-monetary attribute into utility is meant to capture a measure of non-financial plan quality or the impact of constraints such as liquidity constraints, it is not clear why that utility cost would be equivalent in the state of the world where an individual is healthy and does not file any drug claims and the state of the world where she is very ill and files many drug claims. Depending on the context this may or may not be of particular concern, but in this setting, it makes structural interpretation challenging.

This also raises questions regarding how to incorporate those estimates in a counterfactual analysis. Table 4.3 in the main text presents estimates of the conditional logit with and without additional plan attributes. The four specifications imply wildly varying risk aversion of $6.32 \cdot 10^{-4}$, $1.18 \cdot 10^{-3}$, $4.16 \cdot 10^{-6}$, and $-1.08 \cdot 10^{-4}$ (the last value indicating risk loving preferences).

According to the estimates in Column (4) of Table 4.3, a dollar of deductible is equivalent to approximately \$1.18 in premium and \$1.93 in expected out-of-pocket costs. Taking these ratios at face value would suggest that to reduce the deductible from the maximum allowed of \$310 to \$0, a beneficiary is willing to pay approximately \$366 in premiums or \$598 in expected out-of-pocket costs. Such estimates of WTP are obviously challenging to interpret and do not suggest an economic rationale for the estimated importance of the deductible in explaining plan choices. Moreover, the monetary impact of the deductible is already accounted for in the expected out-ofpocket cost. As such, the result that the coefficient on deductible is statistically larger in magnitude than either premium or expected out-of-pocket costs in both Columns (3) and (4), and yet does not correspond to any experienced expenditure, is not consistent with the structural foundation of the model. As shown in the main text in Figure 4.2, the logit regression with additional plan attributes fits the empirical choice probabilities much better than the logit with only cost variables included. Performing a simple model fit Vuong test, however, concludes that this logit model is rejected in favor of the baseline limited consideration model.

A recent alternative explanation for the choice patterns observed in the Medicare Part D market is a model of rational inattention, as presented in Brown and Jeon (2020). In their model, certain attributes of the plans, such as premiums, are easy to observe and known about all available plans. Additional features of the plans, such as OOP costs, are only observed after an individual does research on the plan and pays an information acquisition cost. An implication of their model, discussed as motivating evidence in their study, is a U-shaped relationship between choice quality, measured as the share of people selecting the lowest cost plan, and the stakes of the choice an individual faces. In this case, stakes refer to the variance or standard deviation of costs for an individual across the available plans. In short, for individuals where the stakes are lower and the costs across plans are more similar, you would not expect much research to be conducted as the return to the information cost is limited. For individuals for whom the stakes are high, selecting the wrong plan is costly, and thus, you would expect more research to be conducted. Therefore, the model predicts that individuals with higher stakes will select the optimal plan more frequently than individuals with low stakes. The left-side of the U-shape arises from the fact that for individuals with very low stakes, they face generally low OOP costs and therefore plan selection based predominantly on premiums results in better choice quality than would be true for individuals with higher costs who also do not engage in research. As show in Figure B.2, this reduced form relationship does not manifest in my analysis sample, suggesting a model of rational inattention is not suitable for the choice patterns I observe in this sample.

B.3 Empirical Tests

I present three simple tests of limited consideration within the framework of my model. As descried in more detail in Online Appendix A.2, part of the empirical application in this study involves binning individuals based on similar health and calculating a bin-specific distribution of OOP costs those individuals face under each plan by sampling individuals from the bin and computing their corresponding OOP costs under each plan's specific cost-sharing design. Thus, under an assumption of rational expectations, all individuals in the same health bin who are choosing a plan for the same number of months of 2010 face the same set of distributions of OOP costs across plans, $F(C_{ij}; \theta_{ij}) \forall j \in \mathcal{M}$. Put differently, these individuals differ in their risk preferences, but face the same utility-relevant variables across all plans available.



Figure B.2: Relationship of Stakes and Low Cost Plan Choice

Notes: Stakes are computed as in Brown and Jeon (2020) based on perfect foresight of OOP costs and premiums. Due to sample size, figure plots stakes and share of sample selecting the minimum cost plan at 20 evenly spaced quantiles.

The first test relates to dominant plans. For at least one of these bins of individuals enrolling in plans for the entirety of 2010, at all values of risk aversion a single plan is optimal in terms of expected utility. Therefore, under an assumption of full consideration, all individuals in this bin would select the optimal plan. However, only approximately 26% of individuals within the bin select this dominant plan. A second similar and very simple test evaluates choice shares of non-dominated plans. Such plans are optimal at certain levels of risk aversion, although not at all values of risk aversion. In such a case these plans would be expected under full consideration to have non-zero choice shares, however in two bins such non-dominated plans are never selected.

Finally, a model of full consideration implies the following more precise, third, empirical test.⁴ Consider two bins of individuals b_1 and b_2 , with utility-relevant variables $(\mu_1, \sigma_1^2), (\mu_2, \sigma_2^2)$. Further consider sets $\mathcal{L}_1, \mathcal{L}_2 \subset \mathcal{M}$, and some value of risk aversion $\nu^* \in [0, \bar{\nu}]$. Suppose also that $\operatorname{argmax}_{j \in \mathcal{D}} U(\mu_1, \sigma_1^2; \nu) \in \mathcal{L}_1 \forall \nu \in [0, \nu^*)$ and $\operatorname{argmax}_{j \in \mathcal{D}} U(\mu_1, \sigma_1^2; \nu) \in \mathcal{D} - \mathcal{L}_1 \forall \nu \in (\nu^*, \bar{\nu}]$, and an analogous condition holds for (μ_2, σ_2^2) and \mathcal{L}_2 . Let \hat{j} denote the chosen plan. Then, under the assumption of full consideration, $\operatorname{Pr}(\hat{j} \in \mathcal{L}_1 | (\mu_1, \sigma_1^2)) = \operatorname{Pr}(\hat{j} \in \mathcal{L}_2 | (\mu_2, \sigma_2^2))$. In practice my application

 $^{^{4}}$ See Barseghyan et al. (2021b) for further detail.

presents many possible bins and sets to consider. As an illustration, consider two specific bins: \bar{b}_1 and \bar{b}_2 .⁵ For all levels of risk aversion below $\nu^* = .0033$, the optimal choices for \bar{b}_1 are between two plans $plan_1^1, plan_2^1$ and for \bar{b}_2 two plans $plan_1^2, plan_2^2$. While it is not material, $plan_1^1 = plan_2^2$. For all values of risk aversion above $\nu^* = .0033$, optimal choices are given by $plan_3^1, plan_4^1$ and $plan_3^2$. Given the menu of 46 plans, there are many possible sets \mathcal{L}_1 and \mathcal{L}_2 that satisfy the above condition. The combinatorics make it such that an exhaustive analysis is burdensome, but the test of full consideration is widely rejected in the data, including cases where $Pr(\hat{j} \in \mathcal{L}_1 | (\mu_1, \sigma_1^2)) = .24$ while $Pr(\hat{j} \in \mathcal{L}_2 | (\mu_2, \sigma_2^2)) = .81$. As such, under the specified utility model, full consideration is rejected.

B.4 Additional Geographic Markets

The main analysis in the body of the paper evaluates a single, large geographic market. Due to sample sizes, not all geographic markets in my 5% random sample of Medicare beneficiaries are suitable for estimation. To uncover geographic variation in the choice process, I estimate my baseline model on a number of additional markets in 2010. There are 13 other markets in which I have a sample of at least 2,000 beneficiaries. In two of those markets, New Jersey and Pennsylvania/West Virginia, I am unable to reconcile the observed prices of many drugs in the claims data with the Medicare price files. As such, I am unable to reasonably match the observed OOP drug spending in chosen plans using my cost calculator. To avoid drawing conclusions based on predicted OOP cost distributions that may not reflect the true choice environment of individuals, I only estimate my model in the remaining 11 large markets.

Table B.3 presents results on risk preferences and the main consideration parameters across the 11 regions. In each market, firm effects similarly play a substantial role in the consideration process. Recall the terms δ_{prem} and δ_{ded} reflect the total decay in consideration from the most to least desirable premium and deductible (i.e. $\delta_{ded} = .25$ means the \$310 deductible plan is considered 25% as often as an equivalent \$0 plan).

There is substantial overlap in the confidence intervals for risk preferences, but with some notable exceptions. Specifically, in Illinois and the large midwestern/plains region (bottom row), the model is consistent with beneficiaries exhibiting effectively

⁵Specifically these are two bins where individuals had an average of between two and three monthly claims and were in the second and third quintiles of average drug costs, respectively.

	$Mean(\nu)$	$\operatorname{Var}(\nu)$	δ_{prem}	δ_{ded}
CT/MA/	$[9.41 \cdot 10^{-4}, 2.12 \cdot 10^{-3}]$	$[1.68 \cdot 10^{-6}, 9.93 \cdot 10^{-6}]$	[0.424, 1.000]	[0.208, 0.282]
$\mathrm{RI/VT}$				
NY	$[8.94 \cdot 10^{-4}, 2.09 \cdot 10^{-3}]$	$[1.08 \cdot 10^{-6}, \ 1.02 \cdot 10^{-5}]$	[0.042, 0.080]	[0.276, 0.364]
NG		[2, 0, 0, 1, 0, -7] $(2, 0, 0, 0, 1, 0, -6]$		
NC	$[6.40 \cdot 10^{-4}, 1.59 \cdot 10^{-5}]$	$[3.99 \cdot 10^{-7}, 8.39 \cdot 10^{-6}]$	[0.143, 0.340]	[0.134, 0.192]
FL	$[7.21 \cdot 10^{-4}, 1.33 \cdot 10^{-3}]$	$[2.04 \cdot 10^{-7}, 2.38 \cdot 10^{-6}]$	[0.096, 0.142]	[0.108, 0.145]
	[,]		[,]	[]
AL/TN	$[1.61 \cdot 10^{-4}, 1.08 \cdot 10^{-3}]$	$[1.21 \cdot 10^{-7}, 6.27 \cdot 10^{-6}]$	[0.424, 1.000]	[0.266, 0.350]
MI	$[4.81 \cdot 10^{-4}, 1.50 \cdot 10^{-3}]$	$[1.72 \cdot 10^{-7}, 8.35 \cdot 10^{-6}]$	[0.233, 0.331]	[0.114, 0.159]
ОН	$[6\ 13\ .\ 10^{-4}\ 1\ 51\ .\ 10^{-3}]$	$\begin{bmatrix} 2 & 51 \\ 1 & 10^{-7} \end{bmatrix}$ (2.51 $\begin{bmatrix} 10^{-6} \end{bmatrix}$	[0.053 0.275]	$\begin{bmatrix} 0 & 176 & 0 & 233 \end{bmatrix}$
OII	[0.13 · 10 , 1.51 · 10]	[2.51,10],4.51,10]	[0.000, 0.270]	[0.170, 0.255]
KY/IN	$[1.62 \cdot 10^{-4}, 9.66 \cdot 10^{-4}]$	$[9.92 \cdot 10^{-8}, 5.16 \cdot 10^{-6}]$	[0.035, 0.078]	[0.154, 0.196]
IL	$[2.25 \cdot 10^{-13}, 6.79 \cdot 10^{-4}]$	$[4.13 \cdot 10^{-16}, 3.54 \cdot 10^{-6}]$	[0.494, 1.000]	[0.230, 0.297]
— X				
TX	$[3.23 \cdot 10^{-4}, 9.55 \cdot 10^{-4}]$	$[1.25 \cdot 10^{-7}, 4.25 \cdot 10^{-6}]$	[0.043, 0.114]	[0.185, 0.228]
IA/MN/MT/	$\begin{bmatrix} 1 & 55 \cdot 10^{-13} & 5 & 98 \cdot 10^{-4} \end{bmatrix}$	$[2.07 \cdot 10^{-16} \ 3.46 \cdot 10^{-6}]$	$[0\ 087\ 0\ 174]$	$[0.389 \ 0.467]$
NE/ND/SD/WY	[1.00 10 , 0.00 10]	[=.0, 10 , 0.10 10]		[0.000, 0.101]

Table B.3: Robustness Results: Geographic Variation, Confidence Intervals

Notes: CI based on 1,000 bootstraps. Menus varied from low of 44 in KY/IN to high of 50 in TX. The number of firms offering plans ranged from 18 in KY/IN to 22 in NY.

risk neutrality. Using point estimates rather than confidence intervals, certain patterns about mean risk aversion and local demographics emerge. Using the mean age, share female, and days supply (as a proxy for health) of the samples in each region, I rank the 12 regions I estimate. I find moderate positive correlations between those ranks and the mean of risk aversion.⁶ These results are also suggestive of an intuitive relationship between menus and consideration. In markets with more \$0 deductible plans (by count), the deductible consideration parameter were generally smaller than in markets with fewer \$0 deductible plans (correlation coefficient of -.46). This expected result indicates that as beneficiaries are presented with a greater number of \$0 deductible plans, they are less likely to consider high deductible plans. The same idea can be seen when comparing the deductible parameter to the count of maximal \$310 deductibles (correlation coefficient of .49).⁷ Similarly, the premium consideration effect is stronger (smaller magnitude) in markets with a higher medium premium (correlation of -.68). The results are suggestive that consideration plays a meaningful role across geographies but that the details differ across specific markets.

⁶Specific correlation coefficients are .20, .22, and .29, respectively.

 $^{^{7}}$ If instead one takes the correlation between the \$0 deductible share of the menu and the \$310 share of the menu, the correlation is stronger with -.68 and .56.