A Unified Approach for Ex Ante Policy Evaluation

Achieving (Near) Universal Health Coverage in the US

John Graves, Vanderbilt University School of Medicine

What are the comparative coverage and welfare impacts of expanding health insurance via in-kind public programs versus using tax credits or subsidies to facilitate the purchase of private plans? What is the opportunity cost of making incorrect policy adoption decisions given considerable statistical and modeling uncertainty in these assessments? This study answers these questions by outlining a unified approach for ex ante policy evaluation. By estimating the expected welfare loss from policy adoption decisions based on existing knowledge—which, given current (incomplete) information, may be inconsistent with social preferences–the approach can both inform policy decisions and help refine priorities for future research. An application to US health reform policy demonstrates that based on current information, expansion of in-kind benefits to low-income uninsured populations is optimal provided social preferences are consistent with adopting policies with a ratio of welfare benefits to costs below 0.9. Above this threshold value, a policy that fixes private insurance plan premiums at $5/month for low income populations yields the largest net welfare gain. The information value of reducing parameter uncertainty in these assessments is driven by a small set of parameters that should be prioritized as a focus of future research.

1 Introduction

This study articulates an approach for quantifying the opportunity cost of economic policy decisions made under evidentiary uncertainty. By estimating the expected welfare loss from decisions based on existing knowledge—which, given current (incomplete) information, may be inconsistent with social preferences—the approach can both inform current policy decisions and help refine priorities for future research. An application of these methods examines a central question in U.S. health reform debates: what are the comparative coverage and welfare impacts of expansion of in-kind health insurance programs (e.g., Medicaid) versus using tax credits or subsidies to facilitate the purchase of private insurance plans?

The general setting is one in which empirical and theory-based parameters inform counterfactual projections of policy changes. This type of evaluation—which I will call *ex ante policy evaluation*—is common. Simple (often back-of-the-envelope) counterfactual policy exercies frequently appear at the tail end of research manuscripts. Formal policy models, such as those used by the Congressional Budget Office (CBO), the National Institute of Health Care Excellence (NICE), and other public- and private-sector institutions, are important determinants of the trajectory and outcomes of policymaking worldwide.[1]

Despite the important role played by ex ante evaluation in the policymaking process, scant formal attention has been paid to its integration within the broader economic research enterprise.[2] For example, counterfactual policy evaluations rarely draw on a formal approach to comparative welfare analysis (Hendren and Sprung-Keyser 2020). Rather, models often produce point estimates on an array of welfare-relevant outputs and leave it to policymakers to weigh these separate factors when making decisions (Finkelstein, Hendren, and Shepard 2019; Finkelstein, Hendren, and Luttmer 2019).[3] Moreover, while models often draw on standard economic theory and a shared evidence base, the underlying evidence is estimated with uncertainty—and possibly with bias—and is not always in uniform agreement. Models also differ in their underlying structure, data inputs and assumptions. Finally, the simulation process is often opaque, making it difficult for researchers to understand whether and how their work can inform modeling efforts. Put simply, few economists have a concrete sense of how their research can impact policy modeling done by the CBO and others—nor is it clear whether current research is aligned around questions of highest importance to

[1] Outside of the U.S., entities such as the NICE commission models to determine coverage and reimbursement policy for government-sponsored health programs. Other public- and private-sector institutions have developed simulation models to project changes to tax and transfer policy, education policy, and food policy, among others Avery et al. (2019); Ballas, Clarke, and Wiemers (2006); Basu, Seligman, and Bhattacharya (2013); Cordova et al. (2013); Decoster et al. (2010); Feenberg and Coutts (1993); Kristensen et al. (2014); Mozaffarian et al. (2018); Smith et al. (2006); Sutherland and Figari (2013); Szabó, Gulyás, and Tóth (2008)

[2] The one notable exception is health technology assessment, where rigorous standards for conduct, methods and reporting have been developed by the Panel on Cost-Effectiveness in Health and Medicine (Sanders et al. 2016; Weinstein et al. 1996).

[3] This is particularly true in U.S. health policy, where federal policy decisions based on cost-effectiveness are prohibited through both legislation and administrative rulemaking. But even absent a specific legislative or regulatory decree, the CBO and other modelers have generally avoided producing overall welfare assessments.

informing assessments of the most pressing policy questions.

This study outlines a unified approach that addresses many of these shortcomings. I do so by linking theory for comparative welfare analysis based on the Marginal Value of Public Funds (MVPF; i.e., the ratio of incremental welfare benefits to costs) with approaches for quantifying the welfare loss from policy decisions based on current (incomplete) information (Finkelstein and Hendren 2020; Hendren 2016; Hendren and Sprung-Keyser 2020; Parmigiani and Inoue 2009; Schlaifer and Raiffa 1961).

Intuitively, model uncertainty matters insofar as it affects decisions. Given social preferences and current information, there is a potential welfare loss associated with making the "wrong" decision. For example, at a given policy adoption threshold (e.g., a MVPF value of 0.8, above which a policy might be desirable but below which it may not), parameter uncertainty may or may not affect assessments of whether specific policy changes pass a defined welfare cost-benefit test. If the decision to adopt (or not adopt) a policy is insensitive to parameter variation based on bounds and uncertainty set by current knowledge, then the value of obtaining additional information is low—that is, it is not worth additional effort to reduce information uncertainty since the same policy decision would be made today as it would if had better information. However, if policy adoption decisions are sensitive to uncertainty, then I argue that the MVPF can serve as a lens through which we can quantify the expected welfare loss from policy decisions based on current knowledge. I demonstrate how to estimate that welfare loss, as well as how to decompose it to isolate how specific parameters (or subsets of parameters) contribute to the overall value of information (VOI) across a range of policy adoption thresholds. Parameters with high decision leverage can then be prioritized as a focus of future research.

This study's second contribution is an application of these methods to U.S. health reform policy. Specifically, I calibrate a discrete time and choice model of health insurance take-up. The model itself relies on a relatively circumscribed set of reduced form parameters that capture coverage take-up decisions as well as valuation and costs of in-kind vs. subsidized health insurance benefits (J. Graves et al. 2020; Finkelstein, Hendren, and Shepard 2019; Finkelstein, Hendren, and Luttmer 2019). As such, the underlying model fits within the recent "sufficient statistics" tradition in applied economics (Chetty 2009).

By specifying model parameters in terms of their expected value *and* their underlying uncertainty, I estimate model outcomes based on a sequence of Monte Carlo draws from the joint distribution of parameters. The parameter values and welfare estimates from this exercise are the primary inputs into a "metamodel"—that is, a regression model of how the welfare benefits and costs of specific policy strategies vary as a function of model parameters. Both the theory and results of this study show that metamodels can

facilitate estimation of the value of information based on the expected welfare loss from policy decisions based on current information (Alarid-Escudero et al. 2019; Claxton and Sculpher 2006).

Results from the application yield valuable insights for current policymaking and for future research. If social preferences are consistent with adopting coverage expansions to low-income populations with an MVPF below 0.9, then further expansion of in-kind benefits is optimal given current information. However, for MVPF adoption thresholds above 0.9, a $5/month subsidized private plan minimizes the expected welfare loss as compared with policies to expand in-kind benefits, or that offer less generous subsidies that fix premiums at $25/month or $100/month.

Not surprisingly, there is considerable uncertainty in these assessments. If social preferences are consistent with adopting policies with an MVPF below 0.7, future research should prioritize reducing information uncertainty in parameters summarizing costs to the government among low-income adults when insured by Medicaid and when uninsured. If preferences are more consistent with adopting policies with an MVPF between 0.8 and 1.2, however, research should additionally focus on improving knowledge on the costs and willingness-to-pay among enrollees in private plans, the social marginal utility of income, the pure insurance value of in-kind benefits, and the incidence of uncompensated care. Finally, I extend the modeling framework to explore how additional considerations not yet fully explored in the literature—such as individuals' willingness to pay for medical technology innovation—factor in.

2 Theory

We begin by defining summary measures of policy benefits $W(\boldsymbol{\pi}, \alpha)$ and costs $C(\boldsymbol{\pi}, \alpha)$ as outputs from either an empirical research study or a formal policy simulation model. These measures are defined for each of D total policy strategies $\boldsymbol{\alpha} = (\alpha_1, \cdots, \alpha_D)$ and based on model parameters $\boldsymbol{\pi}$.

Summary measures of benefits and costs are the key ingredients for the Marginal Value of Public Funds (MPVF) developed in a series of studies by Hendren (Hendren and Sprung-Keyser 2020; Hendren 2016). In its most basic form the MVPF is the ratio of policy benefits to costs:

$$MVPF(\boldsymbol{\pi}, \alpha) = \frac{W(\boldsymbol{\pi}, \alpha)}{C(\boldsymbol{\pi}, \alpha)}$$

The MVPF measures the marginal value of an additional dollar spent on a policy. That is, the MVPF quantifies how the welfare benefits accrued by implementing a policy compare to the costs of adopting it. These costs could be mechanical (e.g., the dollar value of a subsidy or cash transfer) and/or the result of

economic frictions brought about through policy implementation (e.g., behavioral changes that result in changes in labor force participation, tax revenue, etc.). The MVPF itself is agnostic: it simply measures the ratio of benefits to costs and does not make affirmative statements about whether a policy is "worth it."

To assess the desirability of policy adoption it is useful to consider a benchmark threshold (λ) summarizing society's willingness to implement a policy. This threshold value could simply be based on a MVPF of 1 or, if society values some redistributive consequence of the policy, could be set based on a value less than one. For example, Finkelstein, Hendren, and Shepard (2019) make comparative assessments of health insurance subsidization policies by specifying a social welfare function over Constant Relative Risk Aversion (CRRA) utility and a defined coefficient of risk aversion ($\sigma = 3$). This results in $\lambda = 0.2$. But researchers do not necessarily have to specify the structure of the social welfare function to define a decision-making benchmark. A value tied to an existing policy with strong social support could also suffice. For instance, Finkelstein, Hendren and Shepard (2017) also consider a benchmark ($\lambda = 0.88$) based on the MVPF of the Earned Income Tax Credit (EITC)—a popular means-tested cash transfer program. Finally, Hendren (2019) argues for the use of efficient welfare weights that project the welfare costs and benefits of any specific policy into the MVPF of a tax transfer between the affected populations.

We can now define the **net welfare benefit**, or the policy's benefits net of its costs scaled by λ

$$NWB(\boldsymbol{\pi}, \alpha, \lambda) = W(\boldsymbol{\pi}, \alpha) - \lambda \cdot C(\boldsymbol{\pi}, \alpha)$$

Intuitively, policies where $NWB \geq 0$ indicate situations where the MVPF is equal to or greater than λ.

2.1 The Opportunity Cost of Imperfect Information

It is important to emphasize that this study does not take a stance on the specific value of λ. Rather, both the theory and results emphasize that the utility of reducing information uncertainty may, for many applications, vary over a range of values for λ. This is readily apparent by noting that the NWB varies as a function of λ.

Relatedly, NWB also depends on model parameters $\boldsymbol{\pi}$, which are often estimated (or assumed) with uncertainty. Uncertainty could derive from sampling or estimation error, or because the parameter values are unexplored in the literature and an educated guess must be made. Alternatively, uncertainty could derive from potential heterogeneity in model parameters across welfare-relevant dimensions (e.g., income, education, geography) or there could be uncertainty around whether literature-based parameters (e.g., elasticities) can reasonably extrapolate to the setting, context and structure of the policy question at hand.

Uncertainty in the "true" value of NWB means that a policy decision from a model based on current knowledge may be inconsistent with social preferences. That is, we might conclude that a policy either does or does not pass a cost-benefit test, but if we had more information on uncertain model parameters we would come to the opposite conclusion. In short, given social preferences, policy adoption decisions based on current information often come with an opportunity cost of making the wrong decision.

To quantify this opportunity cost, we next define the expected welfare gain that would accrue with perfect information on uncertain parameters. Without considering uncertainty, our decision is based on choosing the policy strategy that maximizes NWB with parameters centered on their expected value, i.e., the policy that maps to $\max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi}}[NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)]$. Imagine, however, that we could implement a hypothetical study that eliminates all uncertainty, such that the policy adoption decision is based on $\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)$.

In practice, we may not know parameters with certainty but we can specify their distribution. In that case we can take the expectation over $\boldsymbol{\pi}$ so that with perfect information the expected net welfare benefit is expressed as $E_{\boldsymbol{\pi}}\big[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big]$.

Decision theory defines the difference between these two quantities as the *expected value of perfect information* (EVPI)—that is, it is the expected net welfare change from a hypothetical study that could eliminate all model parameter uncertainty (Schlaifer and Raiffa 1961; Parmigiani and Inoue 2009; Claxton and Sculpher 2006):

$$EVPI(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda) = E_{\boldsymbol{\pi}}\big[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big] - \max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi}}[NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)] \tag{1}$$

It is also useful to think of the EVPI as equivalent to the expected welfare loss (due to information uncertainty) for the optimal policy strategy. To see this, define the optimal strategy as

$$\alpha^* = \arg\max E_{\boldsymbol{\pi}}\big[NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big]$$

so that $NWB(\alpha^*, \boldsymbol{\pi}, \lambda)$ is the net welfare benefit evaluated at the optimal strategy, i.e.,

$$NWB(\alpha^*, \boldsymbol{\pi}, \lambda) = \max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi}}\big[NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big]$$

We can now re-write Equation 1 as

$$EVPI(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda) = E_{\boldsymbol{\pi}}\big[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big] - NWB(\alpha^*, \boldsymbol{\pi}, \lambda) \tag{2}$$

Next, define a welfare loss function

$$L_{\alpha} = \max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda) - NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda) \tag{3}$$

The loss function evaluated at α^* is

$$L_{\alpha^*} = \max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda) - NWB(\alpha^*, \boldsymbol{\pi}, \lambda) \tag{4}$$

And its expected value is given by

$$E_{\boldsymbol{\pi}}\big[L_{\alpha^*}\big] = E_{\boldsymbol{\pi}}\big[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)\big] - NWB(\alpha^*, \boldsymbol{\pi}, \lambda) \tag{5}$$

which is identical to Equation 2.

2.2 Graphical Intuition for the EVPI

Figure 1 provides graphical intuition for the value of information as conveyed through the EVPI. Each scenario plots (as a solid black point) the estimated welfare benefit and cost outcomes from the "baseline" model run (i.e., with all parameters centered at their expected value). The grey points correspond to different model realizations under draws of the parameters from their joint distribution. This "cloud" of points visualizes the degree to which welfare cost and benefit outputs are sensitive to the particular values of parameters used.

As can be seen in the figure, the value of information from reducing parameter uncertainty depends on how we define λ. In each plot, the slope of a ray from the origin has slope $1/\lambda$. Thus, points below each line correspond to combinations of policy benefits and costs that pass the defined cost-benefit test, while those above do not pass the test.

As seen in Scenario A, when λ is low (i.e., λ_1) there little value in obtaining new information: the same decision (to adopt the policy) would be made even if we allowed the model parameters to vary over their entire joint distribution. By comparison, at higher thresholds (λ_2) the value of information is high: the "cloud" of points straddles the threshold line, indicating that different decisions would be made across a series

of parameter draws.

By comparison, in Scenario B there is still variation in modeled costs and benefits but no variation in decisions at either λ_1 or λ_2. Therefore, the value of information is low in both cases. Mathematically, the situation depicted by Scenario B amounts to the loss function having a value of 0 at each realization of our model (and given the choice of λ_1 or λ_2): even under alternative draws of $\boldsymbol{\pi}$, we always choose the same dominant strategy (α^*), so there is no opportunity cost to making the wrong policy adoption decision.

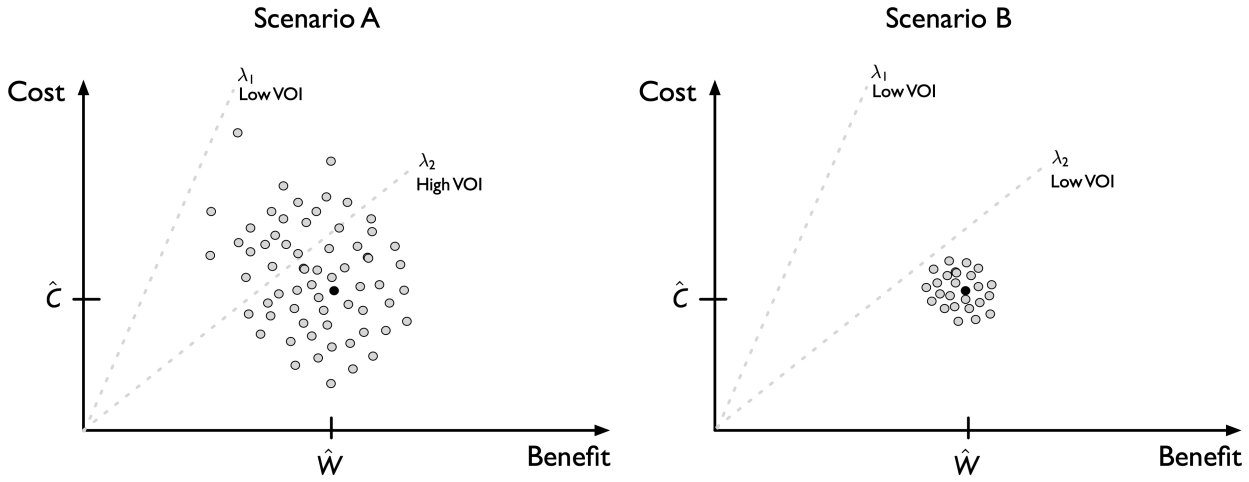

Figure 1: Intuition for Value of Information Theory

2.3 The Value of Information for Parameters of Interest

It is often of interest to isolate the contribution of specific parameters (or sets of parameters) to the overall expected welfare loss from decisions based on current information (Strong, Oakley, and Brennan 2014; Schlaifer and Raiffa 1961; Jalal and Alarid-Escudero 2018). In that case, suppose we could obtain perfect information on the parameter(s) of interest $\boldsymbol{\pi}_z$. The optimal policy decision would then be based on the policy alternative with the highest NWB after averaging over the conditional distribution of remaining parameters $\boldsymbol{\pi}_{-z}$.

$$\max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi}_{-z}|\boldsymbol{\pi}_z} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}, \lambda)$$

However, since $\boldsymbol{\pi}_z$ remains unknown, we we must take the expectation over current information:

$$E_{\boldsymbol{\pi}_z}\big[\max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi}_{-z}|\boldsymbol{\pi}_z} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}_z, \boldsymbol{\pi}_{-z}, \lambda)\big]$$

Echoing equation 1 above, the *expected value of partial perfect information* (EVPPI) is the difference between this quantity and the expected net welfare benefit of the dominant policy strategy (Strong, Oakley, and Brennan 2014; Schlaifer and Raiffa 1961) :

$$EVPPI(\boldsymbol{\pi}_z, \lambda) = E_{\boldsymbol{\pi_z}}\big[\max_{\boldsymbol{\alpha}} E_{\boldsymbol{\pi_{-z}}|\boldsymbol{\pi_z}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}_z, \boldsymbol{\pi}_{-z}, \lambda)\big] - E_{\boldsymbol{\pi}}[NWB(\alpha^*, \boldsymbol{\pi}, \lambda)] \tag{6}$$

2.4 Estimating the Value of Information

In practice, we can estimate the EVPI and the EVPPI by conducting large scale probabilistic sensitivity analyses (Strong, Oakley, and Brennan 2014). To do this, we specify each parameter by its uncertainty distribution and then re-estimate the model K times after taking draws from the joint distribution of parameters. For instance, a model parameter based on a regression coefficient might be specified with a normal distribution centered on the coefficient value and with standard deviation set to the estimated standard error. Alternatively, parameters based on judgment calls might enter the model centered on some value, but with a diffuse (e.g., uniform) prior spread over a plausible range. Finally, it may be the case that parameters co-vary. In that case, we would specify the joint distribution of parameters (e.g., using a multivariate normal distribution centered on estimated regression model coefficients and with variance-covariance matrix as estimated by the same model).

Suppose we construct a K-sized sample of model outputs from a Monte Carlo-based sensitivity exercise. That is, we draw sets of parameters $\boldsymbol{\pi}^{(1)}, \ldots, \boldsymbol{\pi}^{(K)}$ and generate a K-sized vector of net welfare benefits for each policy alternative modeled. As noted above, the EVPI is equivalent to the expected welfare loss for the optimal policy strategy. We can therefore estimate the EVPI by calculating average welfare loss for each strategy in our K-sized probabilistic sensitivity sample:

$$\bar{L}_{\boldsymbol{\alpha}} = \frac{1}{K}\sum_{k=1}^{K}\big[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}, \lambda) - NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}, \lambda)\big] \tag{7}$$

and then taking the minimum across the D strategies.

$$EVPI = \min\{\bar{L}_{\boldsymbol{\alpha}}\} \tag{8}$$

Table 1 walks through a hypothetical example. The table is based on the comparative welfare impacts of two policy options explored in the application below: expansion of health insurance via subsidies ($\alpha = P$) vs. expansion via in-kind benefits ($\alpha = M$). Columns (1) and (2) provide the net welfare benefit (NWB) of

each strategy at a fixed level of λ. Rows correspond to NWB values based on $K = 10$ Monte Carlo draws from the joint distribution of parameters.

As shown in the last row, a strategy of expanding via subsidized private plans yields the highest NWB, on average. This strategy is denoted α^*. Following Equation 7 above, columns (3) and (4) show the welfare loss for each strategy. As shown in Equation 7, the EVPI is equivalent to the strategy that minimizes the mean net welfare loss. Consequentely, the last row of columns (3) and (4) show that the overall value of information is equivalent to 0.074.

Table 1: Example Welfare Loss and Value of Information Calculation

k	Net Welfare Benefit		Welfare Loss	
	NWB_P (1)	NWB_M (2)	$\max(NWB_P, NWB_M) - NWB_P$ (3)	$\max(NWB_P, NWB_M) - NWB_M$ (4)
1	0.478	0.049	0.000	0.429
2	0.256	0.097	0.000	0.159
3	0.486	0.222	0.000	0.265
4	0.221	0.183	0.000	0.038
5	0.232	0.211	0.000	0.021
6	0.454	0.541	0.087	0.000
7	0.221	0.026	0.000	0.194
8	0.271	0.427	0.156	0.000
9	0.353	0.746	0.393	0.000
10	0.225	0.329	0.105	0.000
Mean	**0.320** $= \alpha^*$	0.283	**0.074** = EVPI	0.111

NWB = Net Welfare Benefit; EVPI = Expected Value of Perfect Information

With this hypothetical example in hand, we can next imagine expanding this exercise to include tens of thousands of Monte Carlo draws as well as other values for λ. We can then plot the average welfare loss at each value of λ to construct a set of welfare loss curves—one for each policy strategy under consideration. The lower envelope of this set of curves traces out the strategy, at any given level of λ (and given current information), that minimizes the expected net welfare loss (Alarid-Escudero et al. 2019). As shown in Equation 5 and Table 1, this exercise has the added benefit of providing a VOI estimate at each λ value.

Next, to calculate the EVPPI we can estimate second component of equation 6 as

$$\max_{\boldsymbol{\alpha}} \frac{1}{K} \sum_{k=1}^{K} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}, \lambda) \tag{9}$$

Estimation of the conditional expectation in the first term in equation 6 is less straightforward. Borrowing from the approach in Strong, Oakley and Brennan (2014) we can express model outputs for this term as the sum of the conditional expectation plus a mean-zero error term:

$$NWB(\boldsymbol{\alpha}, \boldsymbol{\pi^{(k)}}, \lambda) = E_{\boldsymbol{\pi_{-z}}|\boldsymbol{\pi_z}=\boldsymbol{\pi_z}^{(k)}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \boldsymbol{\pi}_{-z}, \lambda) + \epsilon^{(k)} \tag{10}$$

In addition, the expectation in equation 10 can be thought of in terms of an unknown function $g(\cdot)$ of $\boldsymbol{\pi_z}$:

$$NWB(\boldsymbol{\alpha}, \boldsymbol{\pi^{(k)}}, \lambda) = g(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \lambda) + \epsilon^{(k)} \tag{11}$$

Based on equation 11, we can estimate the conditional expectation in terms of a "metamodel," or a regression model predicting how the net welfare benefit for a particular policy varies with unknown parameters of interest $\boldsymbol{\pi}_z$. For example, a linear metamodel might specify the function $g(\cdot)$ as a standard regression equation. Alternatively, we might not wish to impose a functional form and instead estimate $g(\cdot)$ nonparametrically. Finally, we might suspect the underlying model has important nonlinearities in the parameters of interest, in which case we can appeal to machine learning methods to estimate $g(\cdot)$. I will explore various approaches to metamodeling in the application below.

3 Application: Modeling Approaches for (Near) Universal Coverage

I next turn to an application of these methods for projecting the coverage and welfare impacts of policies to expand health insurance coverage in the U.S. The policy strategies considered include expansion of in-kind benefits (e.g., Medicaid) vs. the establishment of subsidies to facilitate the purchase of privately-administered plans.

Primary outcomes for this exercise focus on three quantities of interest: counterfactual changes in the distribution of insurance coverage in the US, and comparative welfare impact as measured by the incremental welfare costs and benefits of each policy under consideration.[4] The basic modeling structure is a Markov model of insurance transitions motivated by discrete choice (utility maximization) theory. Additional state values are incorporated into this framework to capture welfare costs and benefits, and to facilitate estimation of the VOI overall and for specific parameters.

While the model is motivated by utility maximization theory, its execution relies on a relatively small set of reduced form parameters. As such, the model fits within the recent "sufficient statistics" tradition in economics (Chetty 2009). Moreover, the modeling framework extends, to the ex ante evaluation setting,

[4]Estimation of the MVPF follows trivially from these quantities, since the MVPF is simply the ratio of benefits to costs.

previous research on the MVPF and VOI of different policies based on ex post evaluation (Hendren and Sprung-Keyser 2020). Finally, because of its simplicity and portability, the model provides a template for researchers who wish to draw on their own reduced form estimates to evaluate counterfactual changes to U.S. health policy without the need for a detailed microsimulation model.[5]

3.1 Discrete Choice Foundations

Consider a model of insurance choice among J alternatives (including the choice not to insure). Define U_{itj} as the utility for choice unit i from selecting choice j at time t.

$$U_{itj} = V(\mathbf{x_{itj}}, \mathbf{z_i}) + \epsilon_{itj} \tag{12}$$

where $\mathbf{x_{itj}}$ is a vector of time-varying attributes of the J choices and the health insurance unit (HIU), or the collection of related family members who could enroll under the same plan. Utility also depends on fixed attributes of the HIU ($\mathbf{z_i}$), and an unobservable component ϵ_{itj}.

For HIU i, the choice of insurance y_{it} is based on maximizing utility across the J alternatives at time t:

$$y_{it} = \arg\max_j [U_{itj}, j = 1, \ldots, J]$$

Next, define a function $B(\cdot)$ mapping utility from choice j to $r_{ij} = P(y_{it} = j)$, the probability of individual i selecting choice j.[6] If the error terms ϵ_{ij} are independent across units and are distributed Type I Extreme Value, we get a standard conditional logit for $B(\cdot)$. However, other link functions—such as based on a nested logit or multinomial logit—could also be used.[7]

3.2 Insurance Choice as a Markovian Process

However we specify the choice probabilities, the choice process at two discrete time periods (t_0, t) can be specified in terms of a Markov trace. Define the *ex ante occupancy vector* $\boldsymbol{p_0}$ summarizing the count or

[5] The Appendix also discusses how the model provides a useful intersection point between reduced form (elasticity-based) modeling approaches common in the applied health economics literature, and formal large-scale microsimulation models as used by the CBO and other research organizations.

[6] As will become important in the next section, we can also think of similar choice probabilities at some time period t conditional on the choice at time $t-1$: $P(y_{it} = s | y_{i,t-1} = r)$.

[7] As discussed in the Appendix, the CBO health reform simulation model uses a nested logit formulation for $B(\cdot)$ in which individuals first select the type of coverage (e.g., employer, non-group, public, uninsured) and then, conditional on that choice, select among available plan types. An alternative approach, however, is to simply specify a reduced form equation that estimates or models $r_{ij} = P(y_{it} = j)$ directly. This type of approach is the basis for elasticity-based microsimulation models previously used by the CBO and used by other modelers (Abraham 2013; Glied and Tilipman 2010; Gruber 2000; "How CBO and JCT Analyze Major Proposals That Would Affect Health Insurance Coverage" 2018; Sonier, Boudreaux, and Blewett 2013).

fraction of the population in each coverage category at time t_0 (i.e., at baseline). Also define a transition probability matrix $\boldsymbol{R_i} = [r_{irs}]$. Cells in this $J \times J$ matrix are defined by transition probabilities based on conditional choice probabilities: $r_{irs} = P(y_{it} = s | y_{i,t_0} = r)$.

With the ex ante occupancy vector and transition probability matrix defined, it is straightforward to obtain the *ex post occupancy vector* summarizing the fraction or number of people in each coverage category at time t:

$$\boldsymbol{p} = \boldsymbol{p_0 R} \tag{13}$$

3.3 Simulating Counterfactual Policy Changes

A key takeaway is that the *ex ante* occupancy vector ($\boldsymbol{p_0}$) and the set of transition probabilities (r_{rs}) can be used to model counterfactual changes to coverage policy. To estimate the value of information and the comparative welfare effects of different policies, we can augment this basic modeling framework by attaching additional state values to coverage transitions. In the application below, for example, I will incorporate additional literature-based parameters summarizing the valuation and take-up of public coverage and subsidized private plans, the costs and incidence of uncompensated care, and social welfare weights that capture social preferences on redistributive transfers of resources to uninsured low-income populations.

To extend the modeling process to explore counterfactual changes, first note that the *ex ante* occupancy vector always remains fixed since it captures the coverage distribution at baseline, i.e., before any counterfactual reform. As such, it can be estimated using nationally-representative survey-based data (e.g., the Current Population Survey, the American Community Survey, etc.). Furthermore, we can express the transition probabilities in potential outcomes notation. That is, define $r_{irs}(0)$ as the probability of transition between coverage categories under the status quo, and $r_{irs}(1)$ as the probability of transition under a counterfactual reform scenario. We can similarly collect these transition probabilities in $J \times J$ transition probability matrices $\boldsymbol{R(0)}$ and $\boldsymbol{R(1)}$.

The impact of reform on the count or fraction of the population in each coverage category is summarized as

$$\begin{aligned} \boldsymbol{\delta} &= \boldsymbol{p(1) - p(0)} \\ &= \boldsymbol{p_0 R(1) - p_0 R(0)} \end{aligned} \tag{14}$$

With this basic modeling framework in hand, we can now efficiently estimate counterfactual changes in the distribution of health insurance; all that is required to model coverage changes associated with a hypothetical reform is estimates of $\boldsymbol{p_0}$ and the transition probabilities. Moreover, as will be discussed in the methods below, the welfare costs and benefits of alternative coverage policies can be estimated by including additional state values attributed to various types of insurance transitions (e.g., uninsured to public vs. uninsured to private).

4 Methods

4.1 Data and Baseline Estimates of Insurance Coverage

Modeling of counterfactual coverage expansion policies begins by specifying and calibrating estimates of insurance transitions under the status quo (i.e., based on $\boldsymbol{p_0}$ and $\boldsymbol{R(0)}$). As covered briefly below and in the Appendix, I estimate the baseline distribution of coverage using the American Community Survey (ACS), and estimate a model of annual coverage changes using panel data on monthly insurance transitions from the Survey of Income and Program Participation (SIPP).

Due to vastly different coverage distributions for individuals below 138% of the poverty line as a result of the 2012 US Supreme Court ruling on the Affordable Care Act's (ACA) Medicaid expansion, I define occupancy vectors separately for expansion vs. non-expansion states (as of 2015). In addition, because individuals can self-report more than one coverage type, I use a coverage hierarchy to classify individuals into their primary source of coverage. This results in the following baseline distribution of coverage:

$$\tilde{\boldsymbol{p}}_{expansion} = \begin{matrix} ESI \\ NG \\ PUB \\ UNIN \end{matrix} \begin{pmatrix} 0.313 \\ 0.082 \\ 0.425 \\ 0.18 \end{pmatrix}$$

$$\tilde{\boldsymbol{p}}_{non-expansion} = \begin{matrix} ESI \\ NG \\ PUB \\ UNIN \end{matrix} \begin{pmatrix} 0.316 \\ 0.104 \\ 0.244 \\ 0.336 \end{pmatrix}$$

where ESI is employer-sponsored insurance, NG is individually-purchased private non-group insurance (e.g.,

ACA Marketplace plans), PUB is public insurance and UNINS is the residual group of individuals without coverage.

I estimate $\boldsymbol{R(0)}$ by fitting a non-parametric (Kaplan-Meier-based) multi-state model of 12-month insurance transitions based on data from Wave 3 of the 2014 Survey of Income and Program Participation (SIPP). The SIPP is a nationally representative panel survey of US households conducted by the US Census Bureau. Beginning with the 2014 panel, SIPP households were interviewed yearly with the goal of evaluating annual and sub-annual changes in insurance coverage, income, and participation in government programs. Each interview covers a look-back period of 12 months. For the 2014 SIPP panel, Wave 3 corresponds to the months covering January-December 2015.

I similarly restrict the SIPP sample to those <138% FPL based on health insurance unit (HIU) income.[8] I also focus on the adult population aged 18 to 63 to avoid age-related changes in insurance coverage attributable to Medicare eligibility at age 65. Finally, I separately define samples for expansion and nonexpansion states (as of 2015).

Twelve-month cumulative transition hazards derived from the multi-state model are the basic building blocks for estimates of the annual coverage transition probabilities contained in $\boldsymbol{R(0)}$. Transition probabilities estimated using the SIPP are then calibrated to match ACS estimates from 2016-2018.[9] Specific details on the multi-state model and Bayesian calibration exercise are provided in the Appendix.

As shown in Figure 2, annual (2016-2018) insurance coverage estimates from the Markov trace (run out to three years) align well with serial cross-section estimates from the ACS. The final calibrated transition probability matrices are also provided below.

$$\boldsymbol{R}(0)_{expansion} = \begin{array}{c} \\ ESI \\ NG \\ PUB \\ UNIN \end{array} \begin{array}{c} \begin{array}{cccc} ESI & NG & PUB & UNIN \end{array} \\ \begin{pmatrix} 0.767 & 0.037 & 0.143 & 0.054 \\ 0.201 & 0.57 & 0.139 & 0.09 \\ 0.102 & 0.019 & 0.838 & 0.041 \\ 0.079 & 0.085 & 0.134 & 0.701 \end{pmatrix} \end{array}$$

[8] An HIU is the collection of related family members who could enroll under the same insurance plan.
[9] Specific details on the multi-state model and calibration process are left to the Appendix.

$$\boldsymbol{R}(0)_{nonexpansion} = \begin{array}{c} \\ ESI \\ NG \\ PUB \\ UNIN \end{array} \begin{array}{c} \begin{array}{cccc} ESI & NG & PUB & UNIN \end{array} \\ \begin{pmatrix} 0.778 & 0.002 & 0.103 & 0.117 \\ 0.27 & 0.53 & 0.154 & 0.046 \\ 0.078 & 0.024 & 0.726 & 0.172 \\ 0.074 & 0.124 & 0.075 & 0.727 \end{pmatrix} \end{array}$$

Uncertainty distributions for the transition probabilities are derived from the posterior distribution output from the calibration exercise, as described in the Appendix.

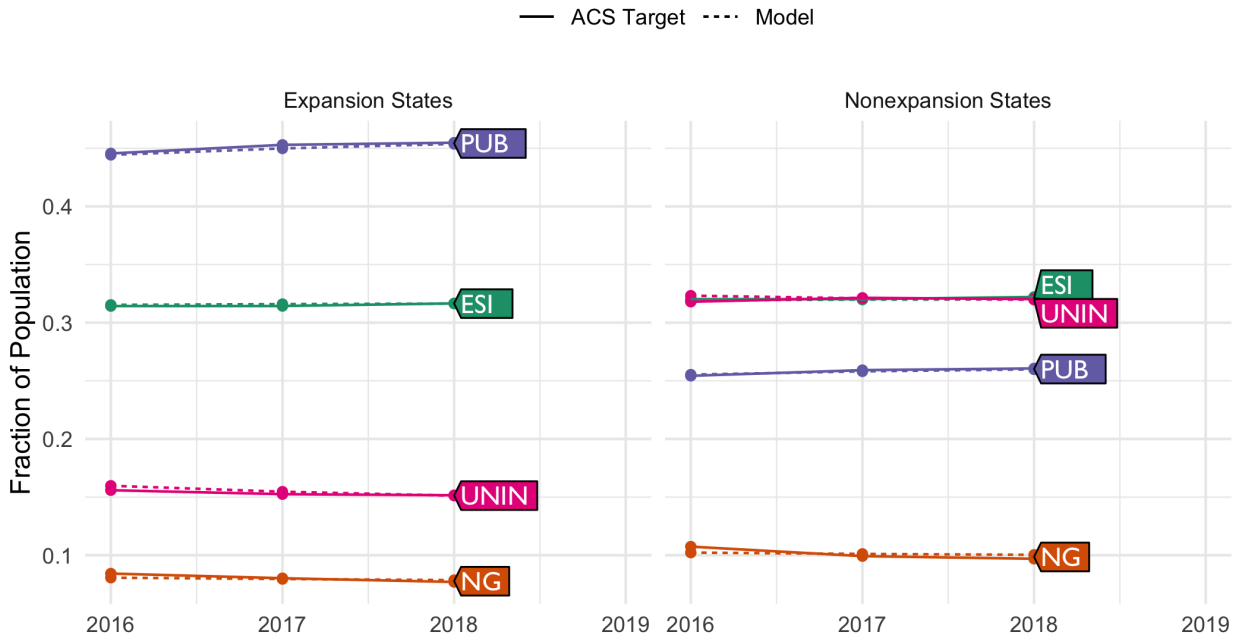

Figure 2: Calibration of Model-Based Estimates to Observed Cross Sectional Insurance Coverage Distribution from the American Community Survey, Adults <138% FPL, 2016-2018.

4.2 Coverage Reform Scenarios

A key feature of the modeling structure developed here is that researchers can simply take the off-the-shelf estimates of $\boldsymbol{p_0}$ and $\boldsymbol{R(0)}$ from above, and then derive their own estimates of $\boldsymbol{R(1)}$ from reduced form evidence (i.e., without the need for a microsimulation model). The application below derives estimates of $\boldsymbol{R(1)}$ and the MVPF for two coverage expansion scenarios. The first is based on evidence on take-up, costs and indivdiuals' valuation of subsidized private plans. The second is based on evidence on takeup, costs and

valuation of in-kind benefits through expansion of a public program option (e.g., Medicaid).

4.2.1 Takeup, Valuation and Costs of Subsidized Private Plans

Estimates of take-up, valuation and costs of expanding coverage via subsidized private plans are derived from Hendren (2020) and Finkelstein, Hendren, and Shepard (2019). These estimates are identified off of a regression discontinuity design that keys off of discrete changes in coverage take-up and enrollee costs at income boundaries set by the subsidy design adopted by Massachusetts' 2006 health reforms.

To model a private plan with generosity comparable to Medicaid (i.e., a plan with little or no cost-sharing), I focus on replicated estimates for the "High Generosity" plan evaluated in Finkelstein, Hendren, and Shepard (2019). I also model takeup of coverage based on scenarios under which beneficiaries pay monthly premiums of $5, $25, or $100—with the remainder of plan costs covered by a federal tax credit or direct subsidy paid to insurers.

These premium scenarios were chosen to closely mirror a set of proposed coverage reform policies to expand the price-linked subsidy design established by the ACA. Under current law, subsidies are only available to households with income between 100% and 400% of poverty. For example, recipients of premium tax credits who have income at the poverty level pay no more than 2.06% of income in monthly premiums and are eligible for additional cost-sharing subsidies that reduce their out of pocket costs to a plan comparable in generosity to Medicaid. For an individual in 2020, the FPL was $12,490. So the $25/month plan modeled here roughly corresponds to the most generous subsidization level under the ACA (2.06% income) being extended to individuals down to 0% FPL.[10]

In addition, the ACA subsidy schedule caps monthly premiums at 9.86% of income up to 400% FPL. Thus, the modeled $100/month plan corresponds (approximately) to a scenario under which the existing income cap is applied to everyone under 400% FPL, but without any additional changes to the subsidy schedule. Finally, under some proposals to reform the ACA, a public option would also be made available to individuals under 138% FPL. The modeled $5 /month premium scenario is intended to capture an option of expanding the ACA tax credits to mimic, as closely as possible, a plan offered in private markets that, like Medicaid, has nominal premium and cost-sharing requirements.

For each scenario, the probability of non-group coverage take-up among the uninsured (i.e., $r_{UNIN,NG}$) is adjusted to reflect takeup at a given subsidized premium level (s_p) based on the willingness-to-pay curve ($D(s_p)$) shown in Figure 3.[11] For example, 90% of the eligible market has a willingness to pay of $15/month

[10] Since current law includes guaranteed issue and community rating for non-group insurance policies, the status quo captures a scenario in which private plans are unsubsidized for low-income individuals.

[11] In principle, additional takeup probabilities could be estimated and incorporated into this framework. For example, if

or above. Therefore, takeup of a price-linked subsidy with a $15/month premium is estimated at 90%.[12]

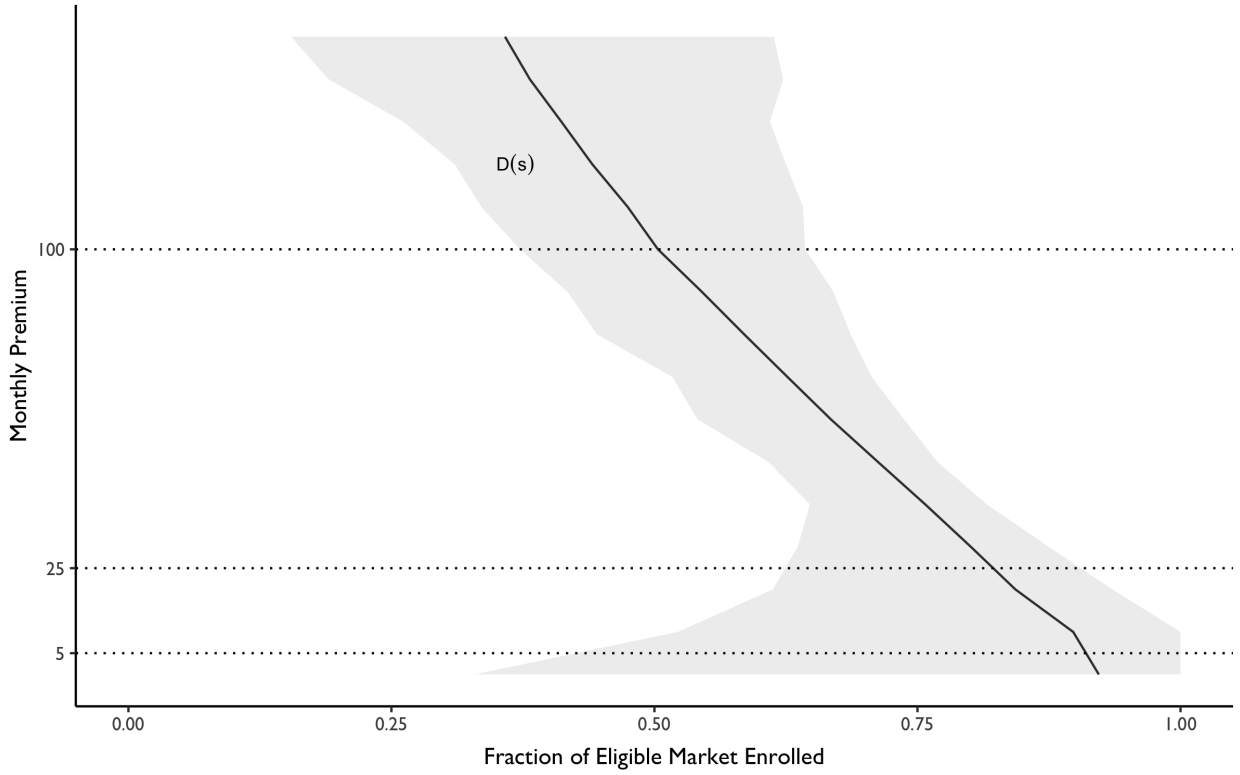

Figure 3: Willingess to Pay Curve with 95% Confidence Interval for Takeup Estimates

As described further below, uncertainty in the probability of take-up is captured by sampling 10,000 draws from a multivariate normal distribution centered on the underlying regression-discontinuity model coefficients and variance-covariance matrix, and then re-deriving $D(s_p)$ based on the process outlined in Finkelstein, Hendren, and Shepard (2019). Ninety-five percent confidence regions for coverage take-up at various WTP values are summarized in the shaded region of Figure 3. This region highlights the high degree of sampling uncertainty in even well-identified evidence of takeup at various subsidization levels—especially for monthly premium levels extrapolated above and below the range of premiums in the underlying Massachusetts premium schedule. This uncertainty will eventually feed into the Monte Carlo-based

generous subsidies for non-group coverage induce firms to drop coverage and send their employees to the marketplace instead, then we would expect additional non-group take-up among people with employer-based insurance. Estimates of welfare costs and benefits could similarly be updated to reflect the welfare costs (e.g., due to "crowd-out" of fully privately-financed coverage) and welfare benefits (e.g., through additional tax revenue as untaxed health insurance benefits are converted into taxable wages as a result of dropping coverage). Such dynamics are beyond the scope of this paper but would be a useful line of inquiry for future research.

[12] A separate but related question is the size of the eligible market. I model this as 50% of the uninsured to capture several features of current law. First, uninsured but unauthorized immigrants are ineligible for tax credits under the ACA. Second, under the ACA families are ineligible for tax credits if a working family member has access to an individual employer-sponsored plan under which the individual contribution is less than about 10% of income (this is the so-called "family glitch" that has restricted eligibility for tax credits). This policy both reduces "erosion" of employer-sponsored insurance (i.e., transitions from employer coverage to subsidized non-group coverage) and further limits the size of the eligible uninsured market.

probabilistic sensitivity exercise described below.

I embed additional parameters within the model to capture the welfare benefits and costs of government subsidies for private insurance plans. These state values are modeled using composite equations based on the the welfare benefit and cost equations in Hendren (2020) and Finkelstein, Hendren, and Shepard (2019).

Welfare benefits for subsidized private plans with out-of-pocket premium p are defined as:

$$W_p = 1 + (1-\theta)\eta \frac{C_U(s_p)}{s_p(-D'(s_p))} + (1-s_p)\zeta(s_p) \tag{15}$$

where s_p is the fraction of the market insured under the scenario, θ is the fraction of uncompensated care borne by the government (e.g., through disproportionate hospital share (DSH) payments and other uncompensated care arrangements), and η is a welfare weight summarizing the social marginal utility of income for other (non-government) parties financing uncompensated care for the uninsured.[13] Equation 15 also captures individuals' willingness to pay for insurance. Specifically, $D'(s_p)$ is the slope of the willingness-to-pay curve at s_p, and $\zeta(s_p)$ captures the additional value of insurance before choices to insure are made (i.e., the ex ante welfare benefit of insurance; see Hendren (2020)).

Intuitively, the first term in Equation 15 captures the observation that (via the envelope theorem) marginal enrollees value an additional \$1 in subsidy at \$1 (see Hendren 2016), the second term captures the additional welfare benefit accruing to those currently bearing some incidence of uncompensated care when those costs are shifted to private plans, and the third term captures the full insurance value against all risks, including those that are observed by plan beneficiaries *before* their choice to insure (Hendren 2020).

Welfare costs for private plans are modeled as

$$C_p = 1 + \frac{C(s_p) - \theta C_U(s_p) - D(s_p)}{s_p(-D'(s_p))} \tag{16}$$

Intuitively, the first term in Equation 16 captures the mechanical cost of an additional \$1 in subsidy, while the second captures lower costs to the government from individuals' premium contributions and from and the reduction in (government) uncompensated care costs when an additional \$1 in subsidy induces additional uninsured individuals to enroll.

As shown in the Appendix, Equations 15 and 16 can be used to replicate—to within a small amount of

[13] This welfare weight can be thought of in terms of the ratio of the social marginal utilities of income to characterize societal willingness to transfer resources from one group (2) to another (1), i.e., $\eta = \frac{\chi_1}{\chi_2}$. Thus, if the incidence of uncompensated care $C_U(s_p)$ among non-governmental parties is on the low-income uninsured themselves, then $\eta = 1$. Alternatively, if the incidence is on higher-income external parties (e.g., providers and hosptials, the privately insured, etc.), then η will be less than 1.

error—the results in Hendren (2020) and Finkelstein, Hendren, and Shepard (2019).

4.2.2 Takeup, Valuation and Costs of Medicaid

Estimates of coverage transition probability changes as a result of public program expansion are drawn from reduced form differences-in-differences (DD) evidence from the ACA's Medicaid expansion (J. Graves et al. 2020). The basis for identification and estimation of changes in transition probabilities stems from a DD design applied to a Markov-based discrete transitions model, and estimated using SIPP data. Details on this model and its assumptions are provided in the Appendix and in J. Graves et al. (2020) and J. A. Graves et al. (2020).

Changes in transition probabilities from implementing an (optional) public program expansion in the remaining nonexpansion states are shown in the DD transition probability change matrix below:

$$\boldsymbol{R}^{DD} = \begin{array}{c} \\ ESI \\ NG \\ PUB \\ UNIN \end{array} \begin{array}{c} \begin{array}{cccc} ESI & NG & PUB & UNIN \end{array} \\ \begin{pmatrix} 0.049 & -0.043 & 0.051 & -0.057 \\ -0.041 & 0.022 & 0.111 & -0.092 \\ -0.014 & -0.008 & -0.018 & 0.04 \\ -0.008 & -0.039 & 0.229 & -0.182 \end{pmatrix} \end{array}$$

Based on this, I model $\boldsymbol{R}(1)$ for nonexpansion states as $\boldsymbol{R}(1) = \boldsymbol{R}(0) + \boldsymbol{R}^{DD}$. For expansion states I do not add in DD estimates since those states already have expanded Medicaid coverage available to those <138% FPL.

I summarize uncertainty around estimates of $\boldsymbol{R}^{DD}$ using a randomization inference procedure as described in J. A. Graves et al. (2020). In short, this procedure re-estimates the underlying differences-in-differences model using a series of 1,000 permuted draws of state expansion status. The collection of K estimates of $\boldsymbol{R}^{DD(k)}$ allows us to approximate a sampling distribution of estimates. Again, details on this process can be found in the Appendix to J. A. Graves et al. (2020).

Parameters capturing indvidiuals' valuation of and the costs of public insurance benefits are derived from the Oregon Health Insurance Experiment (OHIE) and additional assumptions outlined in Finkelstein, Hendren, and Luttmer (2019) (henceforth FHL). To reduce model complexity for this study's application, I do not replicate all the scenarios in FHL, but rather focus on the welfare and cost calculations based on the consumption-based optimization approach. However, the uncertainty ranges specified in the Monte-Carlo

based probabilistic sensitivity exercise cover the relevant ranges of uncertainty captured across the various scenarios modeled in FHL.

Modeled welfare benefit and cost estimates for in-kind public insurance benefits are based on the recognition that coverage provides both a subsidized insurance product to low-income recipients and a monetary transfer to external third parties who subsidize care for the uninsured. Welfare benefits for in-kind programs are summarized as

$$W_m = \theta\gamma + (1-\theta)\big(\gamma + \eta\frac{N}{G}\big) \tag{17}$$

where (as above) θ is the government incidence of uncompensated care and η is the social marginal utility of income, γ is recipients' willingness to pay for Medicaid (defined below), and N is the monetary transfer of Medicaid to external parties. This monetary transfer is defined as the gross cost of Medicaid G minus its net cost (C) to the government, i.e., $N = G - C$. Net costs (C) are defined as

$$C = (G - G_U) + OOP_U$$

where G_U is government spending on Medicaid recipients if they were uninsured, and OOP_U is out-of-pocket spending among these individuals when uninsured. Intuitively, C captures the observation that net costs to the goverment equal gross Medicaid costs minus what the government would spend on recipients if they were uninsured, plus what those recipients would spend while uninsured, since those costs would be covered under the Medicaid benefit.

As noted above, Equation 17 includes γ, which quantifies recipients' willingness to pay for Medicaid benefits. Following FHL, γ is modeled as

$$\gamma = \overbrace{\underbrace{\left[\left(\frac{OOP_U}{G_U}\right) - \left(\frac{OOP_M}{G}\right)\right]}_{\text{Change in OOP price for medical care}} \cdot \underbrace{\frac{1}{2}(G_U + G)}_{\text{Linear approximation}}}^{\text{Transfer Term}} + \overbrace{I}^{\text{Pure insurance term}} \tag{18}$$

where OOP_U and OOP_M are recipients' out of pocket costs when uninsured and insured under Medicaid, respectively. Similarly, G_U and G are government spending on the uninsured and on Medicaid recipients. The pure insurance term, I, in Equation 18 enters the application as a scalar value based on the range of estimates in FHL.[14]

[14] A nice feature of this study's approach is that we can parameterize the model based on a minimal number of parameters.

Finally, the welfare costs of in-kind benefits are represented by

$$C_m = \theta C + (1 - \theta)G \tag{19}$$

where C is the net cost of Medicaid and G is the gross cost of Medicaid, as defined above. In this equation, the first term captures the observation that if the incidence of uncompensated care falls on the government, costs are captured by the net costs after accounting for existing government financing for uncompensated care. By comparison, the second term captures costs when the incidence falls on external (non-government) parties. In that case, costs to the government are the full gross costs.

Again, as shown in the Appendix the welfare cost and benefit equations above successfully replicate the point estimates for the main results in Finkelstein, Hendren, and Luttmer (2019).

4.3 Estimating Coverage Changes and the Expected Welfare Loss

The parameters and equations in the sections above summarize a "sufficient statistics" approach to modeling the coverage and welfare effects of changes to US health insurance policy. This section details an estimation strategy for how these modeled outputs are used to estimate coverage changes and the value of information overall and for individual model parameters.

Changes in the marginal distribution of insurance coverage are modeled using Equation 14 using estimates for $\boldsymbol{p_0}$, $\boldsymbol{R(0)}$ and $\boldsymbol{R(1)}$ as described above. To obtain national estimates, I carry out this exercise separately for nonexpansion and expansion states, and then combine estimates together using total population weights for each state type.

As covered in the theory section, the expected welfare loss for the dominant policy strategy at any given λ is equivalent to the EVPI, i.e., Equation 5. To estimate the expected welfare loss, I carry out probabilistic sensitivity analyses by fitting the model to a sequence of 10,000 draws from the joint distribution of parameters. Each model parameter is assigned an uncertainty distribution based on either sampling uncertainty (e.g., standard errors for point estimates, variance-covariance matrices for parameters derived from regressions, etc.) or based on a uniform prior in cases where the parameter is relatively unexplored in the literature. These uncertainty distributions are summarized in Table 2.

To efficiently carry out this exercise, the 10,000 model runs are based on draws from a Halton

Then, using the VOI methods employed here, we can identify the parameters with the highest amount of leverage. To the extent a parameter such as I proves influential, we can then re-parametrize the model by peeling back the layers of the I onion. In this way we can use the VOI exercise as both a guide for future research on uncertain parameters, and for model building itself.

sequence.[15] This Halton sequence is then mapped (via inverse cumulative distribution functions) to produce a set of 10,000 parameter draws from the distributions specified in Table 2.

[15] A Halton sequence is a sequence of points within [0,1] that appear random but that uniformly cover a multidimensional space. This space could, in principle, be covered using random draws but it would take considerably more draws to ensure uniform coverage across the multidimensional parameter space.

Table 2: Model Parameters

Parameter	Description	Uncertainty Distribution	Source
$\boldsymbol{p_0}$	Baseline distribution of insurance coverage		American Community Survey
r_{rs}	Coverage transition probability r and s		American Community Survey and SIPP
θ	Government incidence of uncompensated care.	$\mathcal{U}(0.5, 1)$	
η	Social marginal utility of income.	$\mathcal{U}(0.25, 1)$	
s	Fraction of eligible market taking up coverage.	$\mathcal{N}(\hat{\boldsymbol{\beta}}, \boldsymbol{\Sigma})$	Finkelstein, Hendren and Shepard (2019)
$D(s)$	Willingness to pay for private insurance	$\mathcal{N}(\hat{\boldsymbol{\beta}}, \boldsymbol{\Sigma})$	Finkelstein, Hendren and Shepard (2019)
$C(s)$	Costs of privately insured	$\mathcal{N}(\hat{\boldsymbol{\beta}}, \boldsymbol{\Sigma})$	Finkelstein, Hendren and Shepard (2019)
C_U	Expected uncompensated care costs if uninsured	$f(OOP_U, G_U, C(s), \phi)$	
OOP_U	Out-of-pocket spending among the uninsured	$\mathcal{N}(569, 100)$	Finkelstein et al. (2012)
G_U	Government spending on the uninsured	$\mathcal{N}(2721, 200)$	Finkelstein et al. (2012)
ϕ	Moral hazard effect of insurance.	$\mathcal{N}(0.25, 0.05)$	Finkelstein et al. (2012)
$\zeta(s)$	Difference in marginal utilities between insured and uninsured	$f(D(s), \sigma)]$	Hendren (2020)
σ	Risk aversion parameter	$5e-4$, $\mathcal{U}(0, 1e-3)$	Handel et al. (2015); Hendren (2020)
γ	Willingness to pay for Medicaid benefits	$f(Tr, I)$	Finkelstein, Hendren and Shepard (2019)
Tr	Transfer term	$f(OOP_U, G_U, OOP_M, G)$	Finkelstein, Hendren and Shepard (2019)
G	Gross Medicaid spending	$\mathcal{N}(3600, 200)$	Finkelstein et al. (2012)
OOP_M	Medicaid recipients' out-of-pocket spending	0	Finkelstein et al. (2012)
I	Pure insurance term	$\mathcal{N}(760, 100)$	Finkelstein, Hendren and Luttmer (2019)
N	Monetary transfer of Medicaid to external parties	$f(G, C)$	Finkelstein, Hendren and Luttmer (2019)
C	Net cost of Medicaid to the government	$f(G, G_U, OOP_U)$	Finkelstein, Hendren and Luttmer (2019)

I next execute the underlying Markov-based model on each of the 10,000 parameter set draws. Mean values for changes in each coverage category, as well as welfare cost and benefit outcomes, are the basis for "baseline" estimates (i.e., estimates with each parameter centered on its expected value), while 95% intervals are constructed with upper and lower limits set at the 2.5th and 97.5th quantiles, respectively.

Estimates of the overall value of information are calculated using the set of estimates from the sensitivity exercise, as well as a sequence of values for λ ranging from 0.2 to 1.6. Specifically, using the welfare cost and benefit outcomes, I estimate $\hat{\alpha}^*_\lambda$ as the strategy that maximizes the average NWB at a given λ value:

$$\hat{\alpha}^*_\lambda = \arg\max \left[\frac{1}{K} \sum_{k=1}^{K} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}, \lambda) \right]$$

The average welfare loss for the dominant policy strategy provides an estimate of the EVPI:

$$EVPI_\lambda = \bar{L}_{\alpha^*,\lambda} = \frac{1}{K} \sum_{k=1}^{K} \left[\max_{\boldsymbol{\alpha}} NWB(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}, \lambda) - NWB(\hat{\alpha}^*_\lambda, \boldsymbol{\pi}^{(k)}, \lambda) \right]$$

As noted in the theory section, the set of welfare loss estimates for each policy strategy can also be plotted with varying values for λ to trace out a set of welfare loss curves. The lower envelope of this curve provides an estimate of the EVPI and the dominant policy strategy (i.e, the strategy that minimizes the expected welfare loss) at each value of λ.

4.4 Metamodeling to Estimate the Value of Information for Specific Parameters

The EPVI provides an estimate of whether it is worth pursuing additional research to reduce information uncertainty. However, a key downside is that this value assumes *all* uncertainty is eliminated; it does not provide guidance on which parameters, or subsets of parameters, contribute the most (or least). Understanding how the EVPI decomposes into indivdual model inputs is important so that future research can focus on questions of paramount importance to improving policy decisions based on better information.

To estimate the VOI for individual parameters and subsets of parameters I draw on metamodeling techniques, as discussed in the theory section above. Specifically, the collection of parameter draws, as well as the welfare cost and benefit outcomes from each model run, are the basic inputs into a metamodel for estimating Equation 11:

$$NWB(\boldsymbol{\alpha}, \boldsymbol{\pi^{(k)}}, \lambda) = g(\boldsymbol{\alpha}, \boldsymbol{\pi}^{(k)}_z, \lambda) + \epsilon^{(k)}$$

Because the function $g(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \lambda)$ has an unknown form, I execute metamodels using two strategies: (1) linear regression with a flexible (third-order) polynomial for the parameter(s) of interest; and (2) using a generalized additive model (GAM).[16] A GAM represents the function $g(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \lambda)$ using a series of smooth functions of each input.

$$g(\boldsymbol{\alpha}, \boldsymbol{\pi}_z, \lambda) = s_1(\pi_1) + ... + s_p(\pi_p)$$

The results below utilize a smoothing function based on a cubic spline with three knots.

Outcomes for the metamodel are the $K \times D$ matrix of net welfare benefits for each policy strategy from the probabilistic sensitivity exercise. I next cycle through a series of metmodels fit to predictors which could be individual model parameters or sets of model parameters that (collectively) capture key domains of the simulation model.

Based on output from the metamodel I estimate the EVPPI as:

$$\widehat{EVPPI}(\boldsymbol{\alpha}, \boldsymbol{\pi}_z, \lambda) = \frac{1}{K} \sum_{k=1}^{K} \left[\max_{\boldsymbol{\alpha}} \hat{g}(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \lambda) \right] - \max_{\boldsymbol{\alpha}} \frac{1}{K} \sum_{k=1}^{K} \hat{g}(\boldsymbol{\alpha}, \boldsymbol{\pi}_z^{(k)}, \lambda) \tag{20}$$

5 Results

Table 3 summarizes the main results. Based on current information, a strategy to expand in-kind Medicaid benefits to low-income populations would result in a 4.6 percentage point decline in the national uninsured rate among adults <138% FPL (approximately 3 million people), and carries an MVPF of 0.92 (95% confidence interval 0.64 to 1.39). By comparison, offering heavily subsidized $5/month private insurance plans to low income populations would reduce the uninsured rate by 5.9 percentage points, and has an MVPF of 0.90 (95% confidence interval 0.71 to 1.26). Policies that fix subsidies at $25 or $100 per month would cover fewer uninsured people and yield MVPF estimates of 0.81 and 0.85, respectively.

[16] Given the nonlinearity of many policy simulation models, other approaches such as a gaussian process regression or a machine learning model could also suffice. For the application here I found broad agreement in results across alternative approaches including ordinary least squares (polynomial) regression, GAMs, and artificial neural networks, so only focus on two metamodeling strategies in the main results.

Table 3: Coverage and Welfare Estimates

		Coverage Expansion Scenarios (ppt. change)			
	Baseline	Medicaid	$5/month	$25/month	$100/month
Coverage Distribution Changes					
Employer-Sponsored	30.0	1.1 [-0.6,2.8]	0.1 [-0.0,0.2]	0.1 [0.0,0.1]	0.0 [-0.0,0.1]
Private Non-Group	8.7	-1.0 [-2.5,0.5]	5.7 [0.8,7.2]	4.7 [2.5,6.1]	1.8 [0.4,3.3]
Public	36.5	4.6 [3.0,6.3]	0.1 [-0.0,0.2]	0.1 [0.0,0.2]	0.0 [-0.0,0.1]
Uninsured	24.8	-4.8 [-6.8,-2.8]	-5.9 [-7.6,-0.8]	-4.9 [-6.4,-2.6]	-1.8 [-3.4,-0.3]
Welfare Impact					
Welfare Benefit		1.30 [0.77,2.45]	1.05 [0.93,1.34]	1.11 [1.01,1.34]	1.26 [1.03,1.67]
Welfare Cost		1.41 [1.02,2.14]	1.12 [0.76,1.65]	1.39 [1.13,1.68]	1.52 [0.95,1.97]
Marginal Value of Public Funds		0.92 [0.64,1.39]	0.90 [0.71,1.26]	0.81 [0.69,0.94]	0.85 [0.63,1.19]

95% confidence intervals shown in brackets.
ppt. = percentage point

As the results above make clear, the comparative coverage and welfare effects of expanding in-kind benefits vs. subsidies are quite similar—but there is considerable uncertainty in these estimates. This uncertainty is depicted in Figure 4. This figure is similar in structure to Figure 1. However, instead of plotting a "cloud" of 10,000 points, the figure plots 95% ellipses for each expansion strategy.[17]

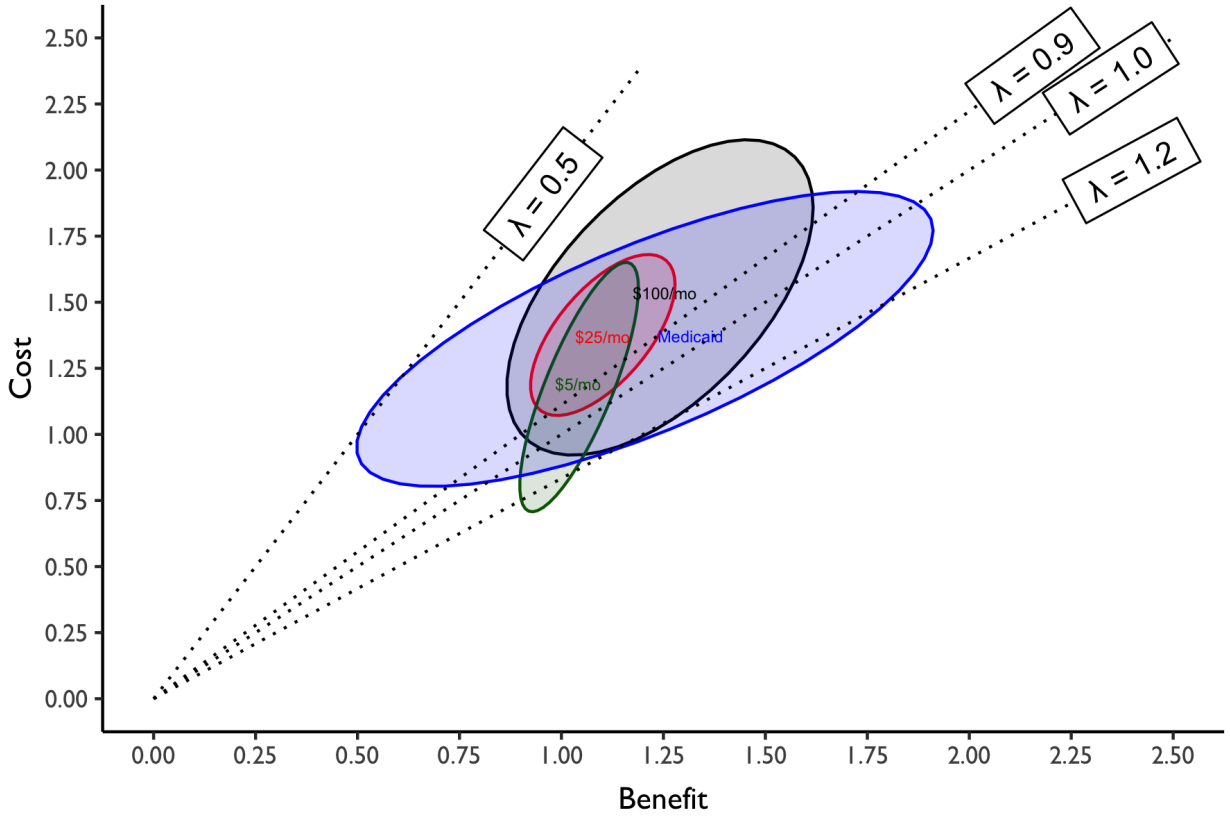

Figure 4: Cost and Benefit Estimates

As Figure 4 demonstrates, uncertainty in the welfare benefit and cost estimates for all strategies straddle meaningful policy adoption threshold values (i.e., different values of λ). In addition, the shape of each ellipse is slightly different. For the \$5/month strategy there is more sensitivity to parameter uncertainty in estimates of welfare costs than welfare benefits.[18] However, for the Medicaid strategy there is far more uncertainty in welfare benefits, with values ranging from 0.5 to nearly 2.0.

Figure 5 plots the average welfare loss curves for each strategy. As mentioned in the theory and methods sections above, the lower envelope of this set of curves provides both an estimate of the optimal

[17] Each ellipse covers approximately 95% of estimates from the sensitivity exercise.

[18] This observation makes intuitive sense given that welfare cost estimates (i.e., Equation 16) depend heavily on the estimated cost of marginal enrollees in private plans. As shown in the Appendix and the in the partial VOI results below, there is considerable uncertainty in these costs, especially for premium regions outside the region captured by the Massachusetts exchange subsidy schedule.

policy strategy (i.e., the strategy that maximizes the NWB, given current information) at any given λ value, as well as an estimate of the overall value of information as measured using the EVPI.

Based on this frontier, I find that given current information, a strategy to expand in-kind benefits is optimal provided social preferences are consistent with an MVPF policy adoption threshold of around 0.94 or below. For adoption thresholds above 0.94, a strategy to expand by fixing premiums at \$5/month for low-income populations yields the highest NWB and lowest expected welfare loss from imperfect information.

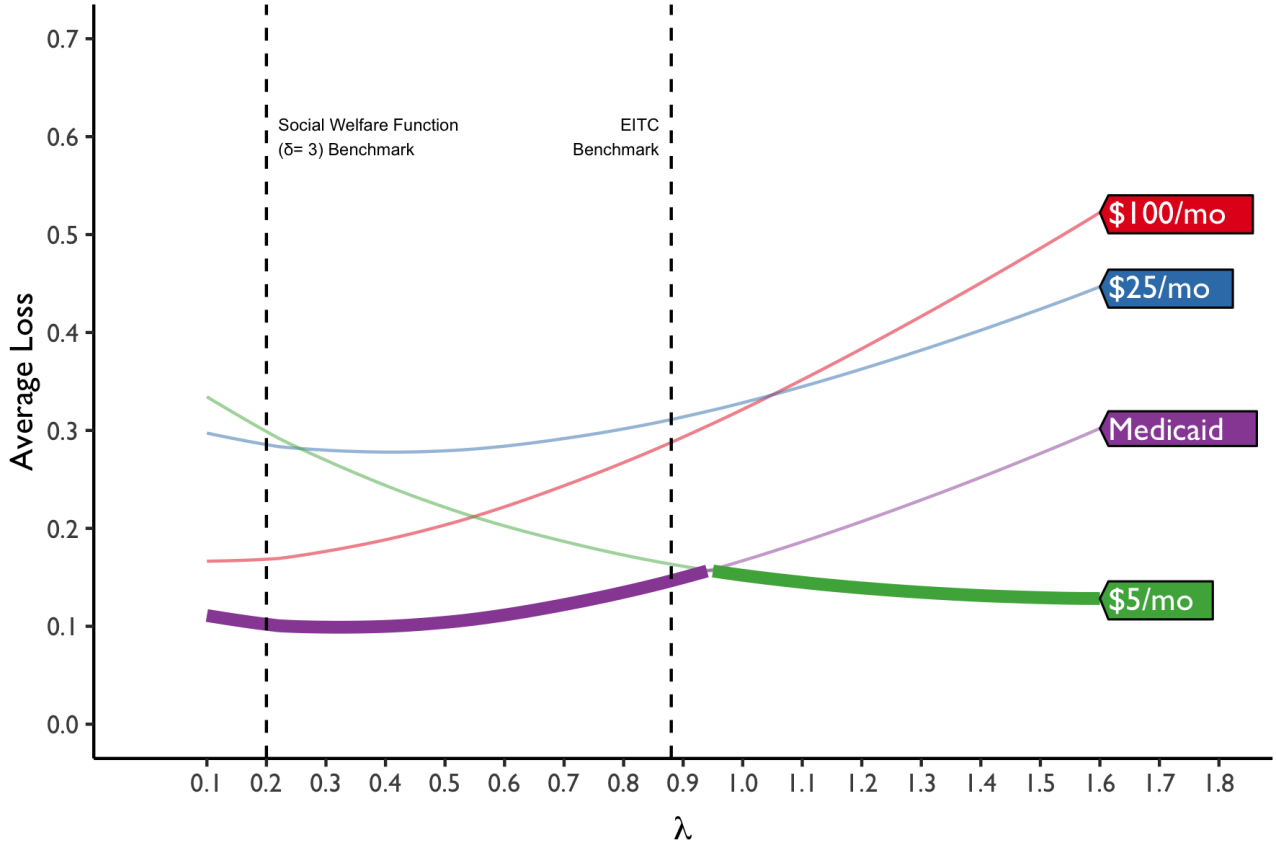

Figure 5: Expected Loss Curves

To provide some context around these estimates, Figure 5 also plots (as vertical lines) reference thresholds based on the implied λ value for a social welfare function based on constant relative risk aversion ($\sigma = 0.3$; CRRA) individual utility (as explored in Finkelstein, Hendren, and Shepard (2017)), as well as a λ benchmark based on the EITC as estimated in Hendren (2016). Interestingly, the "switching point" between optimal strategies (0.94) is close to the EITC MVPF. Given this finding, and given that the EITC is a popular means-tested cash transfer program (i.e., it appears to pass an implicit societal cost-benefit test) it is perhaps not surprising that debates over universal coverage in the U.S. remain unsettled over the question of

whether to expand via public benefits vs. via government-subsidized private coverage—and, moreover, that different states (e.g., Arkansas) have elected to use subsidized private insurance benefits in lieu of Medicaid even under the ACA.

Figure 6 provides EVPPI estimates as derived from metamodels. To construct this figure, the overall parameter set is partitioned into subsets that correspond to input parameters as summarized in Table 2. I then fit a series of metamodels to each of these parameter subsets and across a variety of values for λ. The EVPPI value from each of these models is then plotted. The x-axis covers a range of values for λ while the y-axis plots the EVPPI for each parameter set at any given value of λ. Each panel corresponds to a different model specification, either based on a linear ordinary least squares (with 3rd degree polynomial) regression, or a GAM.

For values of λ below 0.5, the parameters with the highest information value correspond to government costs among Medicaid recipients (G), and to the government costs for those recipients when uninsured (G_U). Both of these parameters are based on experimental evidence from the OHIE, and the difference between them is the causal effect of Medicaid on government spending (i.e., $G - G_U$). This OHIE estimate is based on a local average treatment effect of Medicaid of approximately \$780, and that has a relatively high standard error (\$371).[19] It is this sampling uncertainty, and its importance in determining the valuation of Medicaid benefits relative to its costs, that informs the (relatively high) estimate of partial VOI shown in Figure 6.

Both G and G_U retain high information value for MVPF adoption thresholds approaching 1.0. However, for λ values above 0.7, additional parameters summarizing the costs and takeup of private coverage contribute meaningfully to the overall value of information. Sampling uncertainty in takeup of private plans at different premium levels was apparent and noted in Figure 3.[20] This statistical uncertainty is important for determining the welfare costs and benefits associated with expanding coverage using subsidized private plans.

While statistical uncertainty from existing experimental and reduced form evidence explains a large fraction of the VOI, additional uncertainty attached to the values of other structural components of the model also contribute—but only at policy adoption thresholds in the neighborhood of 0.9. For example, the pure insurance value of Medicaid (I), the social marginal utility of income (η) and the incidence of uncompensated care (θ) all have relatively high EVPPI values. By comparsion, parameters summarizing the moral hazard effects of insurance (ϕ), and risk aversion parameters used to determine the ex ante value of insurance (σ) play an important role in the model, but do not play an important role in shifting decisions among the policy alternatives considered in the application.

[19] See Table V in Finkelstein et al. (2012)

[20] An additional figure showing statistical uncertainty around estimates of costs can be found in the Appendix.

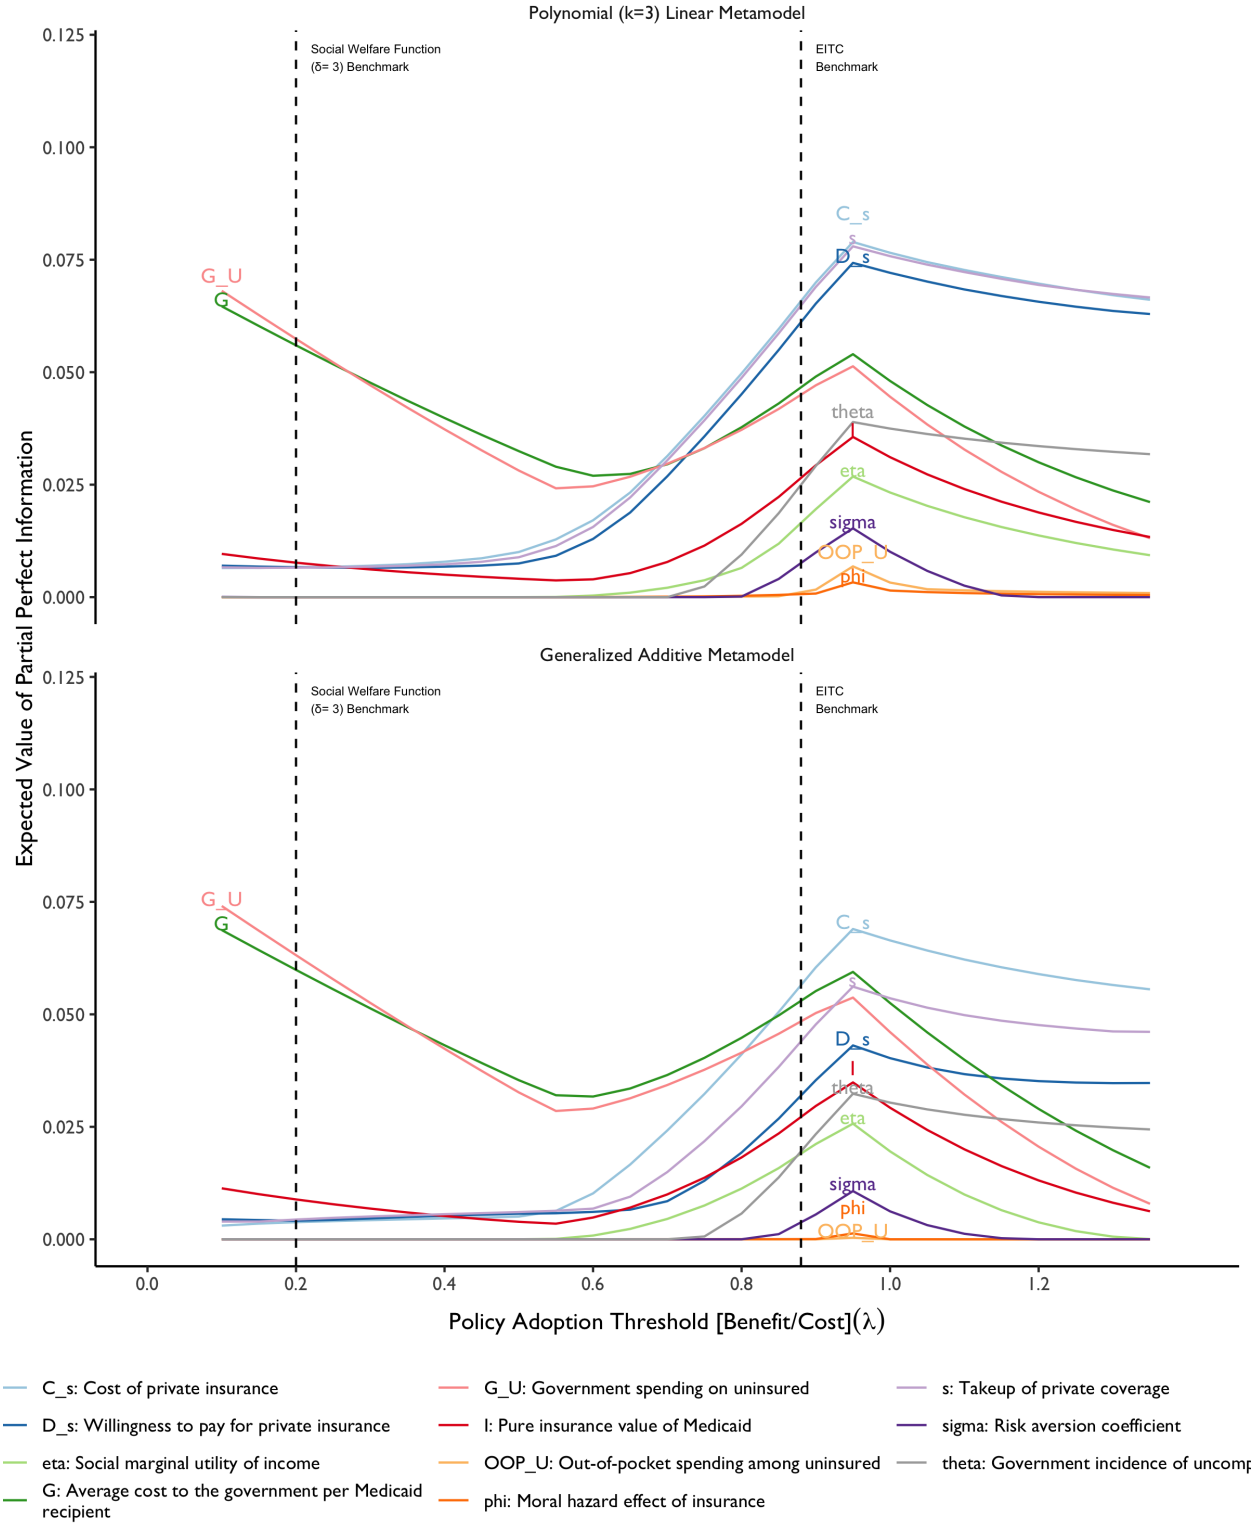

Figure 6: Metamodel Estimates of the Contribution of Individual Parameters to the Overall Value of Information

6 Model Extensions: Welfare Impact of Health Technology Innovation

It is worth emphasizing that while the application above draws heavily on the published literature, the overall approach is not dependent on either the structure or specific parameter values derived from previous research. Rather, this study's methods can extend to any number of structural or reduced form methods for quantifying the incremental welfare costs and benefits of policy changes.

Similarly, readers may find that other important welfare determinants may be missing from the model's structure. This "structural uncertainty" goes beyond the statistical uncertainty explored above, and would therefore not be captured in the VOI estimates in the previous section. As this section will show, however, an additional benefit of this study's approach is that the underlying infrastructure can also be used to guide model building itself.

This section augments the modeling application to consider an additional input into welfare cost-benefit considerations: the market-altering impact of coverage expansions on incentives for medical technology innovation. This consideration may be an important but unexplored dimension because previous large-scale insurance expansions (e.g., introduction of the Medicare program in 1965) have been shown to shift the incentive landscape in health care markets (Finkelstein 2007). Thus, further expansion of either in-kind or private health insurance could further shift either (or both) the scope and profile of technologic innovation. In particular, it is an open and important question whether large-scale expansion of in-kind public programs with administratively-set prices may trade off access to future medical technology (which may never be developed if the price incentives are weaker) for broader access to today's medical technology (Chandra and Garthwaite 2019).

We will operationalize these considerations by including an additional parameter (ι) that captures individuals' willingness to pay for the option value of future medical technology. This option value arises if a predominantly private health insurance system differentially increases future innovation by creating a more fertile development & pricing environment relative to a largely public insurance-based system.

Because this parameter is unexplored in the literature, we attach a relatively uniform prior to its underlying distribution, and allow it to vary between 0 and 0.04. In terms of magnitude this corresponds, roughly, to up to twice the magnitude of individuals' willingness to pay for learned health risk information before the decision to insure is made (i.e., β; see Hendren 2020). Thus, these are relatively large bounds that could be improved by focused future research on this topic. The ι parameter is then added to the welfare

benefit equation for private plans, and is also added to the welfare cost of in-kind expansions.

Figure 7 below reproduces Figures 5 and 6 above, but with the innovation parameter included. With innovation effects the frontier shifts, resulting in a break-point at MVPF thresholds of 0.75, as compared with 0.94 in the main results. In other words, to the extent individuals highly value the option value of future medical technologies that would be developed with more widespread private (vs. public) coverage, social preferences would need to accommodate policies with an MVPF of 0.75 or below for a strategy of in-kind program expansion to be optimal given current information.

Again, however, the specific value of ι is relatively unknown because it has not yet been explored in the literature. Nevertheless, this concept remains important, as evidenced by the EVPPI curves shown in the bottom panel of Figure 7. As a consequence, in addition to research on the parameters identified above, further research on the welfare consequences of innovation under large-scale expansions would seem both important and warranted.

7 Conclusion

Uncertainty is inherent to all models, but matters insofar as it affects policy adoption decisions. This study drew on recent advancements in public finance and decision theory to formalize and embed this intuition within counterfactual policy projection models. The application developed a novel sufficient statistics approach to modeling the coverage and welfare impacts of changes to US health insurance policy. The results thus speak to the costs and benefits of expanding coverage based on existing knowledge, but also show how researchers can leverage underlying model uncertainty to provide a guidepost for the direction of future research.

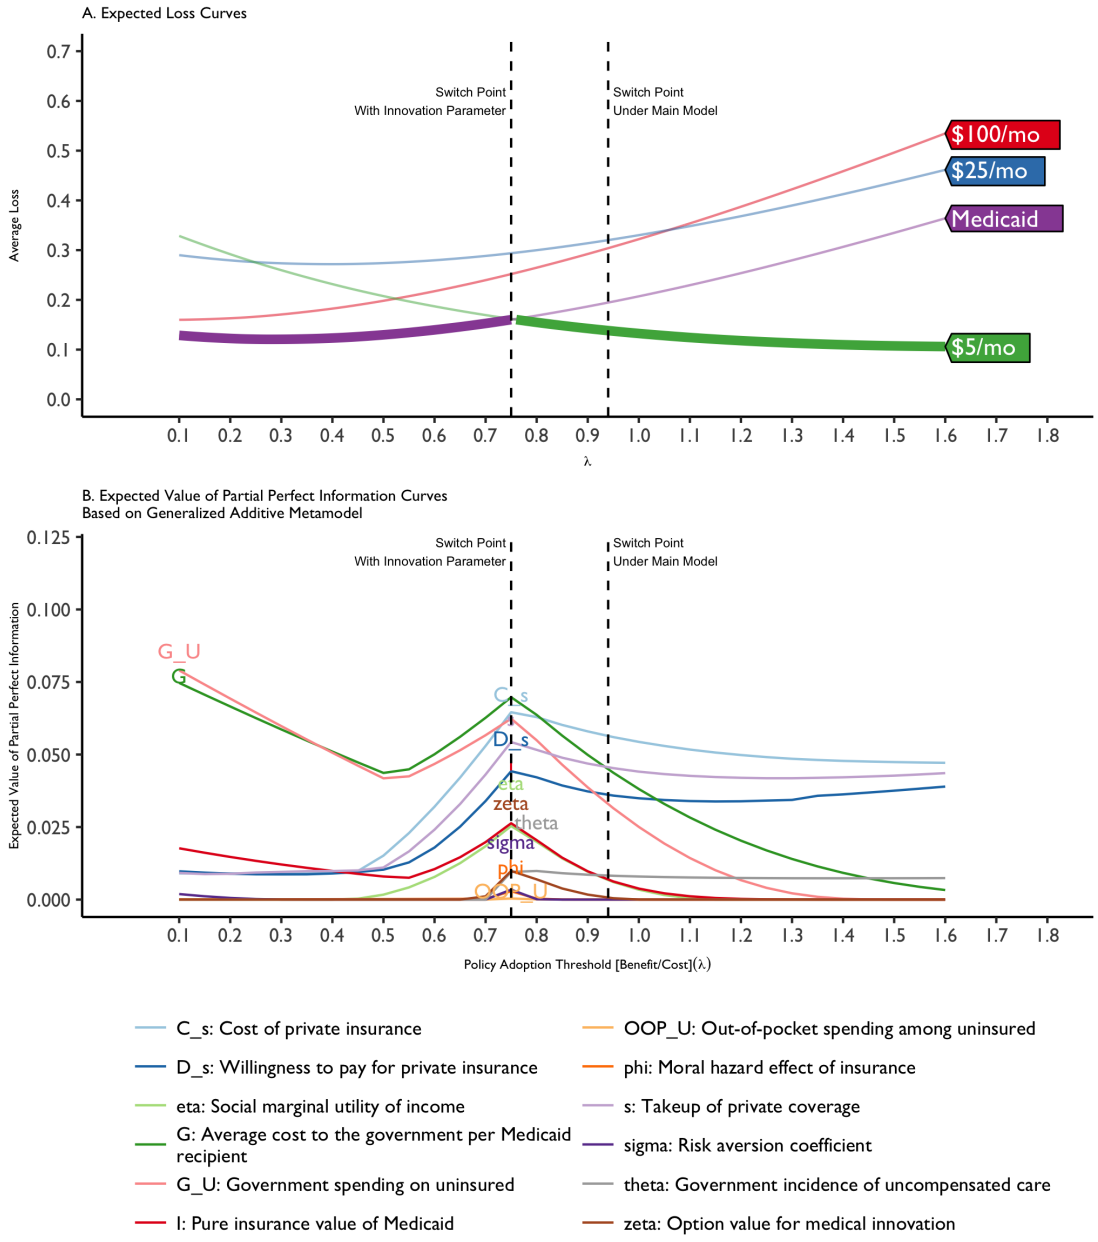

Figure 7: Expected Loss and Expected Value of Partial Perfect Information Curves with Innovation Effects Included

8 References

Abraham, Jean Marie. 2013. "Using Microsimulation Models to Inform US Health Policy Making." *Health Services Research* 48 (2pt2): 686–95.

Alarid-Escudero, Fernando, Eva A. Enns, Karen M. Kuntz, Tzeyu L. Michaud, and Hawre Jalal. 2019. "'Time Traveling Is Just Too Dangerous' but Some Methods Are Worth Revisiting: The Advantages of Expected Loss Curves Over Cost-Effectiveness Acceptability Curves and Frontier." *Value in Health* 22 (5): 611–18. https://doi.org/10.1016/j.jval.2019.02.008.

Avery, Christopher, Jessica Howell, Matea Pender, and Bruce Sacerdote. 2019. "Policies and Payoffs to Addressing America's College Graduation Deficit." *Brookings Papers on Economic Activity*, no. Fall 2019 (September).

Ballas, D., G. P. Clarke, and E. Wiemers. 2006. "Spatial Microsimulation for Rural Policy Analysis in Ireland: The Implications of CAP Reforms for the National Spatial Strategy." *Journal of Rural Studies* 22 (3): 367–78.

Basu, Sanjay, Hilary Seligman, and Jay Bhattacharya. 2013. "Nutritional Policy Changes in the Supplemental Nutrition Assistance Program: A Microsimulation and Cost-Effectiveness Analysis." *Medical Decision Making* 33 (7): 937–48.

Chandra, Amitabh, and Craig Garthwaite. 2019. "Economic Principles for Medicare Reform , Economic Principles for Medicare Reform." *The ANNALS of the American Academy of Political and Social Science* 686 (1): 63–92. https://doi.org/10.1177/0002716219885582.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annu. Rev. Econ.* 1 (1): 451–88.

Claxton, Karl P., and Mark J. Sculpher. 2006. "Using Value of Information Analysis to Prioritise Health Research: Some Lessons from Recent UK Experience." *PharmacoEconomics* 24 (11): 1055–68. https://doi.org/10.2165/00019053-200624110-00003.

Cordova, Amado, Federico Girosi, Sarah Nowak, Christine Eibner, and Kenneth Finegold. 2013. "The COMPARE Microsimulation Model and the US Affordable Care Act." *International Journal of Microsimulation* 6 (3): 78–117.

Decoster, André, Jason Loughrey, Cathal O'Donoghue, and Dirk Verwerft. 2010. "How Regressive Are Indirect Taxes? A Microsimulation Analysis for Five European Countries." *Journal of Policy Analysis and Management* 29 (2): 326–50.

Feenberg, Daniel, and Elisabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management* 12 (1): 189–94.

Finkelstein, Amy. 2007. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare." *The Quarterly Journal of Economics* 122 (1): 1–37.

Finkelstein, Amy, and Nathaniel Hendren. 2020. "Welfare Analysis Meets Causal Inference." *Journal of Economic Perspectives* 34 (4): 146–67.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6): 2836–74. https://doi.org/10.1086/702238.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2017. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." National Bureau of Economic Research.

———. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530–67.

Glied, Sherry, and Nicholas Tilipman. 2010. "Simulation Modeling of Health Care Policy." *Annual Review of Public Health* 31: 439–55.

Graves, John A., Laura A. Hatfield, William Blot, Nancy L. Keating, and J. Michael McWilliams. 2020. "Medicaid Expansion Slowed Rates Of Health Decline For Low-Income Adults In Southern States: An Analysis of the Impact of Medicaid Expansion on the Self-Reported Health of Low-Income Older Nonelderly Adults Living in the South." *Health Affairs* 39 (1): 67–76.

Graves, John, Carrie Fry, Michael McWilliams, and Laura A Hatfield. 2020. "Difference-in-Difference Estimation for Transitions in Discrete Outcomes: Insurance Transitions After the Affordable Care Act's Medicaid Expansion." Working Paper.

Gruber, Jonathan. 2000. "Microsimulation Estimates of the Effects of Tax Subsidies for Health Insurance." *National Tax Journal*, 329–42.

Hendren, Nathaniel. 2016. "The Policy Elasticity." *Tax Policy and the Economy* 30 (1): 51–89.

———. 2020. "Measuring Ex Ante Welfare in Insurance Markets." *The Review of Economic Studies*. https://doi.org/10.1093/restud/rdaa015.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3): 1209–318.

"How CBO and JCT Analyze Major Proposals That Would Affect Health Insurance Coverage." 2018. Congressional Budget Office.

Jalal, Hawre, and Fernando Alarid-Escudero. 2018. "A Gaussian Approximation Approach for Value of Information Analysis." *Medical Decision Making* 38 (2): 174–88.

Kristensen, Alyson H., Thomas J. Flottemesch, Michael V. Maciosek, Jennifer Jenson, Gillian Barclay, Marice Ashe, Eduardo J. Sanchez, Mary Story, Steven M. Teutsch, and Ross C. Brownson. 2014. "Reducing Childhood Obesity Through US Federal Policy: A Microsimulation Analysis." *American Journal of Preventive Medicine* 47 (5): 604–12.

Mozaffarian, Dariush, Junxiu Liu, Stephen Sy, Yue Huang, Colin Rehm, Yujin Lee, Parke Wilde, Shafika Abrahams-Gessel, Thiago de Souza Veiga Jardim, and Tom Gaziano. 2018. "Cost-Effectiveness of Financial Incentives and Disincentives for Improving Food Purchases and Health Through the US Supplemental Nutrition Assistance Program (SNAP): A Microsimulation Study." *PLoS Medicine* 15 (10): e1002661.

Parmigiani, Giovanni, and Lurdes Inoue. 2009. *Decision Theory: Principles and Approaches*. Vol. 812. John Wiley & Sons.

Sanders, Gillian D., Peter J. Neumann, Anirban Basu, Dan W. Brock, David Feeny, Murray Krahn, Karen M. Kuntz, David O. Meltzer, Douglas K. Owens, and Lisa A. Prosser. 2016. "Recommendations for Conduct, Methodological Practices, and Reporting of Cost-Effectiveness Analyses: Second Panel on Cost-Effectiveness in Health and Medicine." *Jama* 316 (10): 1093–103.

Schlaifer, Robert, and Howard Raiffa. 1961. *Applied Statistical Decision Theory*.

Smith, Dianna M., Graham P. Clarke, Joan Ransley, and Janet Cade. 2006. "Food Access and Health: A Microsimulation Framework for Analysis." *Studies in Regional Science* 35 (4): 909–27.

Sonier, Julie, Michel H. Boudreaux, and Lynn A. Blewett. 2013. "Medicaid 'welcome-Mat'effect of Affordable Care Act Implementation Could Be Substantial." *Health Affairs* 32 (7): 1319–25.

Strong, Mark, Jeremy E. Oakley, and Alan Brennan. 2014. "Estimating Multiparameter Partial Expected Value of Perfect Information from a Probabilistic Sensitivity Analysis Sample: A Nonparametric Regression Approach." *Medical Decision Making: An International Journal of the Society for Medical Decision Making* 34 (3): 311–26. https://doi.org/10.1177/0272989X13505910.

Sutherland, Holly, and Francesco Figari. 2013. "EUROMOD: The European Union Tax-Benefit Microsimulation Model." *International Journal of Microsimulation* 6 (1): 4–26.

Szabó, Attila, László Gulyás, and István J. Tóth. 2008. "TAXSIM Agent Based Tax Evasion Simulator." In *5th Conference of The European Social Simulation Association, University of Brescia, Italy.*

Weinstein, Milton C., Joanna E. Siegel, Marthe R. Gold, Mark S. Kamlet, and Louise B. Russell. 1996. "Recommendations of the Panel on Cost-Effectiveness in Health and Medicine." *Jama* 276 (15): 1253–58.

Supplemental Appendix: A Unified Approach for Ex Ante Policy Evaluation

Evaluating Mechanisms for (Near) Universal Health Coverage

John Graves, Vanderbilt University

Contents

1 Modeling Coverage Changes

Recall from the main text that the impact of a modeled reform on coverage is summarized as

$$\boldsymbol{\theta} = \boldsymbol{p(1)} - \boldsymbol{p(0)} = \tilde{\boldsymbol{p}}'\boldsymbol{R(1)} - \tilde{\boldsymbol{p}}'\boldsymbol{R(0)} \tag{1}$$

where $\tilde{\boldsymbol{p}}$ is the *ex ante occupancy vector* summarizing the count or fraction of the population in each coverage category at time $t-1$ (i.e., at baseline). The transition probability matrix is defined as $\boldsymbol{R_i} = [r_{irs}]$. Cells in this $J \times J$ matrix are defined by transition probabilities among the J possible coverage categories based on conditional choice probabilities: $r_{irs} = P(y_{it} = s | y_{i,t-1} = r)$.

1.1 Estimating $\tilde{\boldsymbol{p}}$

The ex ante occupancy vector summarizes the fraction or count of the target population in each coverage category at baseline. We model coverage changes based on the baseline distribution of insurance coverage in 2015 as estimated using the American Community Survey (ACS). Due to vastly different coverage distributions for individuals below 150% of the poverty line as a result of the 2012 US Supreme Court ruling on the ACA's Medicaid expansion, we define occupancy vectors separately for expansion vs. non-expansion states (as of 2018). In addition, because individuals can self-report more than one coverage type, we use a coverage hierarchy to classify individuals into their primary source of coverage. This results in the following baseline distribution of coverage:

$$\tilde{\boldsymbol{p}}_{expansion} = \begin{pmatrix} \texttt{ESI} = 0.313 \\ \texttt{NG} = 0.082 \\ \texttt{PUB} = 0.425 \\ \texttt{UNIN} = 0.18 \end{pmatrix}$$

$$\tilde{\boldsymbol{p}}_{non-expansion} = \begin{pmatrix} \texttt{ESI} = 0.316 \\ \texttt{NG} = 0.104 \\ \texttt{PUB} = 0.244 \\ \texttt{UNINS} = 0.336 \end{pmatrix}$$

1.2 Estimating and Calibrating $\boldsymbol{R(0)}$

As a prerequisite to modeling counterfactual coverage changes under a hypothetical reform, we first estimate and calibrate the baseline (status quo) transition probability matrix $\boldsymbol{R(0)}$. This allows us to specify an initial Markov trace that, when modeled forward over time (i.e., $\tilde{\boldsymbol{p}}'\boldsymbol{R(0)}^{\boldsymbol{T}}$, where T is the number of years to model forward), yields estimates of the marginal distribution of insurance coverage that closely match the actual distribution observed in Census survey data in each year from 2015-2018. We do so by first fitting a non-parametric multi-state model of insurance transitions based on Waves 2 and 3 of the 2014 Survey of Income and Program Participation (SIPP).

The SIPP is a nationally representative panel survey of US households conducted by the US Census Bureau. Beginning with the 2014 panel, SIPP households were interviewed yearly with the goal of evaluating annual and sub-annual changes in insurance coverage, income, and participation in government programs. Each

interview covers a look-back period of 12 months. For the 2014 SIPP panel, Wave 2 corresponds to January-December 2014, and Wave 3 corresponds to January-December 2015.

We restrict the SIPP sample to those <150% FPL based on health insurance unit (HIU) income.[1] We also focus on the adult population aged 18 to 63 to avoid age-related changes in insurance coverage attributable to Medicare eligibility at age 65. Finally, we separately define samples for expansion and nonexpansion states (as of 2015).

Health insurance (including uninsured status) is measured for each month an individual is in sample using the primary source of coverage hierarchy. We then identify transitions among primary coverage source and fit a nonparametric (Kaplan-Meier-based) multi-state model to estimate transition hazards. See Graves and Nikpay (2017) for a detailed explanation of multi-state modeling of insurance changes.

Twelve-month cumulative transition hazards derived from the multi-state model are the basic building blocks for estimates of annual coverage transition probabilities. Specifically, we collect 12-month cumulative transition hazards into a rate matrix $\boldsymbol{G}$, with diagonal elements set to 1 minus the sum of other row elements. The estimated rate matrices for expansion and nonexpansion states are shown below.

$$\boldsymbol{G}(0)_{expansion} = \begin{pmatrix} -0.082 & 0.037 & 0.026 & 0.019 \\ 0.125 & -0.186 & 0.04 & 0.021 \\ 0.01 & 0.006 & -0.036 & 0.02 \\ 0.032 & 0.02 & 0.175 & -0.227 \end{pmatrix}$$

$$\boldsymbol{G}(0)_{non-expansion} = \begin{pmatrix} -0.083 & 0.033 & 0.014 & 0.036 \\ 0.082 & -0.161 & 0.04 & 0.039 \\ 0.018 & 0.014 & -0.079 & 0.047 \\ 0.028 & 0.024 & 0.072 & -0.123 \end{pmatrix}$$

The rate matrix facilitates embedding of a transition probability matrix based on a yearly time step.[2] This is accomplished by taking the matrix exponential of $\boldsymbol{G}$:

$$\boldsymbol{R(0)}^{init} = e^{\boldsymbol{G}}$$

The initial transition probability matrices are listed below.

$$\boldsymbol{R}(0)^{init}_{expansion} = \begin{pmatrix} 0.924 & 0.033 & 0.027 & 0.017 \\ 0.11 & 0.833 & 0.039 & 0.018 \\ 0.01 & 0.006 & 0.966 & 0.018 \\ 0.029 & 0.017 & 0.155 & 0.799 \end{pmatrix}$$

[1] An HIU is the collection of related family members who could enroll under the same insurance plan.

[2] Any time step could be chosen, however we elect to model an annual time step since most health insurance plans are annual contracts.

$$\boldsymbol{R}(0)^{init}_{non-expansion} = \begin{pmatrix} 0.922 & 0.03 & 0.015 & 0.033 \\ 0.073 & 0.853 & 0.037 & 0.036 \\ 0.018 & 0.013 & 0.926 & 0.043 \\ 0.027 & 0.022 & 0.066 & 0.887 \end{pmatrix}$$

For a detailed explanation of embedding transition probability matrices based on a rate matrix see the Appendix to J. Graves, Garbett, et al. (2020).

1.2.1 Calibration to American Community Survey Data

While in principle we could use the embedded transition matrix above as our primary estimate of $\boldsymbol{R(0)}$, doing so can result in diverging estimates (relative to Census data) of the marginal coverage distribution over time. For example, running a Markov trace forward from 2015 (based on $\tilde{\boldsymbol{p}}$ and $\boldsymbol{R}(0)$ as defined above) results in marginal coverage distribution estimates for 2018 that differ slightly from cross-sectional 2018 estimates from the ACS. This is readily apparent by comparing the "Uncalibrated" (purple) and "ACS Target" (Green) lines in Figure 1, which shows the estimated fraction of the (expansion state) population in each category from 2015-2018.

In part this divergence may be due to secular trends that affect the likelihood of coverage transition. To the extent these secular changes reflect a process not in equilibrium, this can manifest in length-biased spells in the SIPP data due to left-truncation (Asgharian, M'Lan, and Wolfson 2002; de Uña-Alvarez, Otero-Giráldez, and Alvarez-Llorente 2003). In other words, the "stock" sample of individuals in each coverage category at the beginning of 2014 may be overly representative of individuals in longer coverage spells.

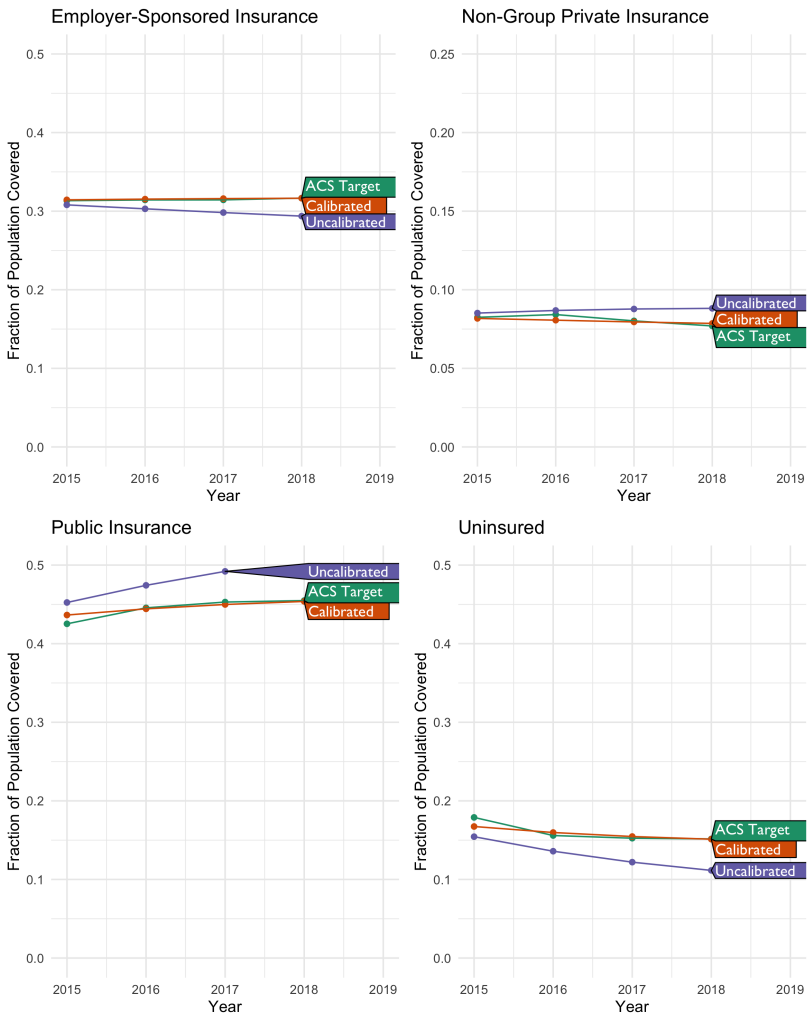

Figure 1: Calibration Plot for Distribution of Insurance Coverage by Year

To address this, we implement a calibration procedure to adjust transition probabilities so that modeled coverage transitions under the status quo match calibration targets set by the ACS. We do so by jointly calibrating transition probabilities using incremental mixture approximate Bayesian computation (Raftery and Bao 2010; Rutter et al. 2019). Approximate Bayesian computation is an efficient method for simulating draws from the posterior distribution of model parameters given a set of calibration targets. In our case, priors are set using the initial embedded transition matrix $\boldsymbol{R}(\mathbf{0})^{init}$, and calibration targets are derived from yearly (2015-2018) estimates from the ACS.

Densities for the posterior distribution of each transition probability are shown in Figures 2 and 3.

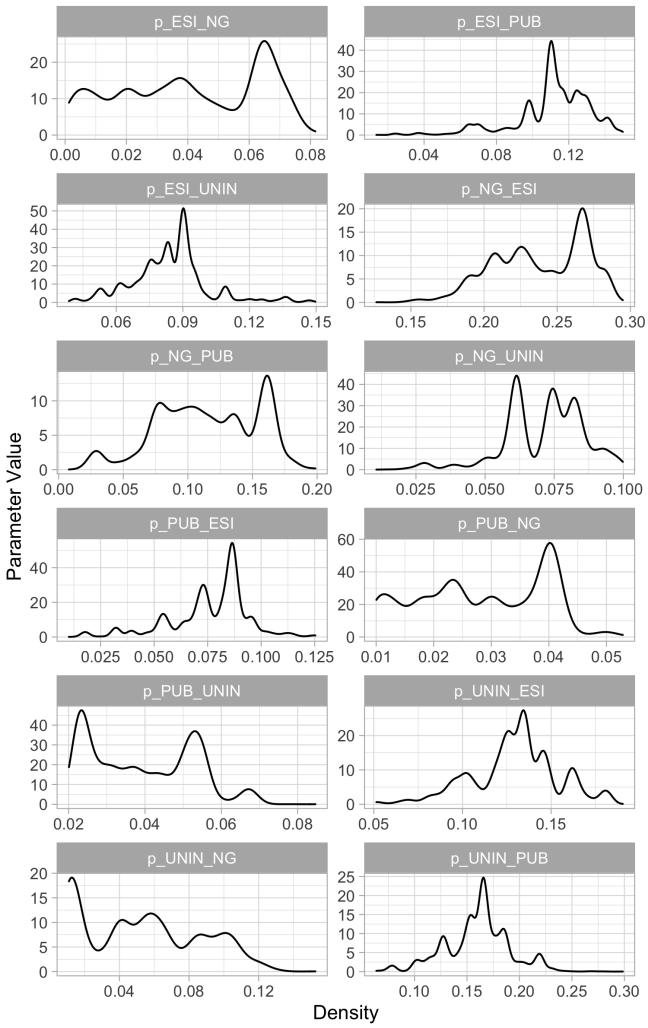

Figure 2: Expansion States: Posterior Distribution of Calibrated Transition Probabilities

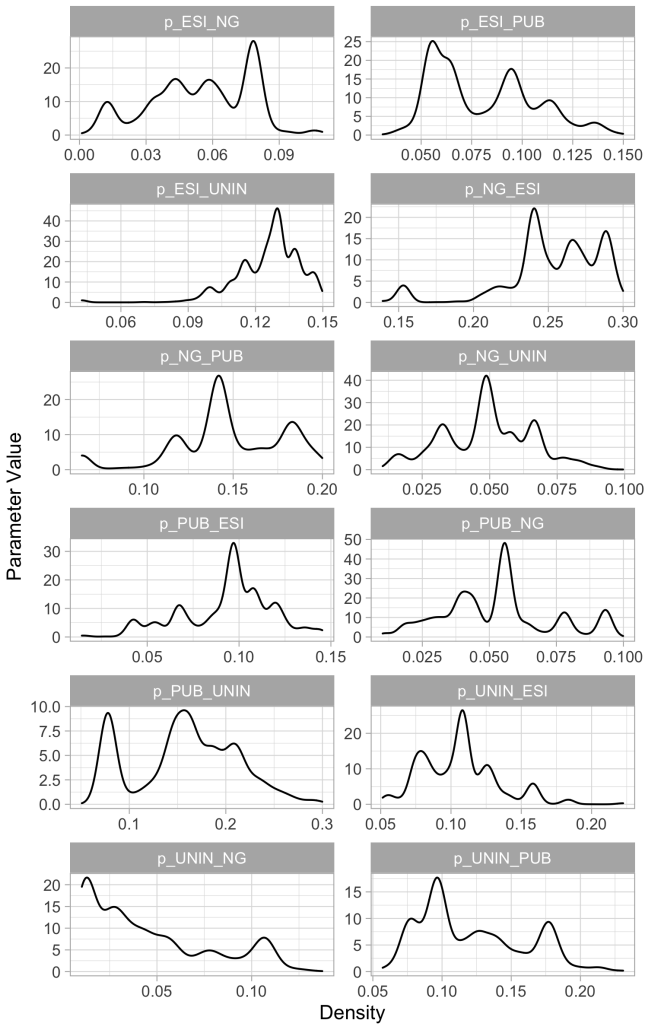

Figure 3: Non-Expansion States: Posterior Distribution of Calibrated Transition Probabilities

After calibration we obtain the following baseline transition probability matrices for expansion and nonexpansion states, respectively.

$$\boldsymbol{R}(0)_{expansion} = \begin{pmatrix} 0.767 & 0.037 & 0.143 & 0.054 \\ 0.201 & 0.57 & 0.139 & 0.09 \\ 0.102 & 0.019 & 0.838 & 0.041 \\ 0.079 & 0.085 & 0.134 & 0.701 \end{pmatrix}$$

$$\boldsymbol{R}(0)_{non-expansion} = \begin{pmatrix} 0.778 & 0.002 & 0.103 & 0.117 \\ 0.27 & 0.53 & 0.154 & 0.046 \\ 0.078 & 0.024 & 0.726 & 0.172 \\ 0.074 & 0.124 & 0.075 & 0.727 \end{pmatrix}$$

Consistent with the hypothesis that the SIPP's "stock" sample of spells in progress in January 2015 may be

disproportionately representative of longer spells, we find that the calibrated transition probability matrix results in higher probabilities for transitions *out of* coverage types—and correspondingly lower probabilities for remaining in the same coverage type after one year.

Modeling forward estimates of $\tilde{\boldsymbol{p}}$ (for 2015) and $\boldsymbol{R}(0)$ yield marginal coverage distributions that match observed coverage distributions as estimated by the ACS. This can be seen by comparing the "Calibrated" (orange) and "ACS Target" lines in Figures 1 and 4.

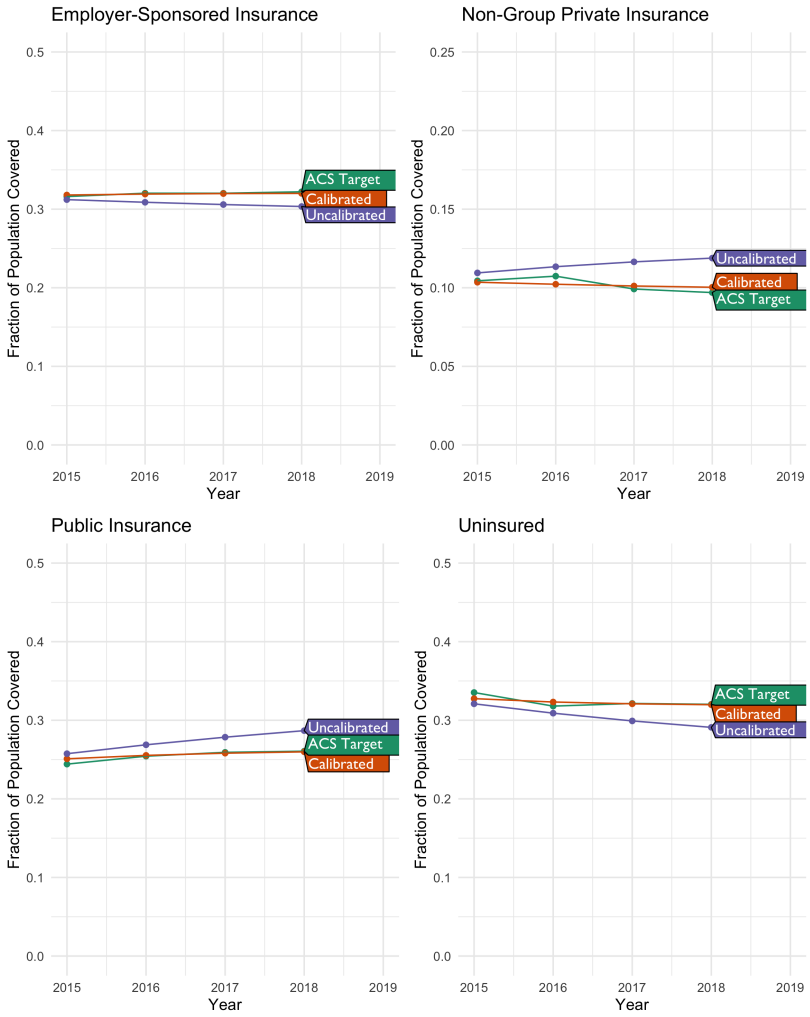

Figure 4: Non-Expansion States: Calibration Plot for Distribution of Insurance Coverage by Year

1.3 Modeling $\boldsymbol{R}(1)$: Link to CBO and Other Microsimulation Models

With estimates of $\tilde{\boldsymbol{p}}$ and $\boldsymbol{R}(0)$ in hand, we now have a set of parameters that, when modeled forward in a Markov trace, result in marginal coverage distribution estimates that closely match cross-sectional ACS estimates from 2015 to 2018. This set of estimates constitutes a scenario under the status quo, or with no further changes to US coverage policy.

To model counterfactual changes, we next turn to estimates of $\boldsymbol{R(1)}$ under our two hypothetical reform

scenarios: (1) further expansion of public insurance to those <150% FPL in nonexpansion states; (2) universal expansion of subsidized private insurance to all individuals <150% FPL.

Estimates of $\boldsymbol{R(1)}$ can be obtained in a number of ways. One option is simply to curate the literature for reduced form evidence that can inform calculating or deriving changes in coverage transition probabilities under a modeled reform. Another option is to model changes in choice probabilities directly using a microsimulation model.

While our application relies on the former (literature-based) approach, we first discuss briefly how to derive choice probabilities from either a utility-maximization-based or elacticity-based microsimulation model. This discussion is useful because it highlights how the overall modeling structure summarized in Equation 1 provides a useful intersection point between diverse methods for ex ante evaluation used by researchers, the CBO, and others.

A standard assumption is that an exogenous policy change does not affect the unobserved disturbance term ϵ_{itj} in an underlying discrete choice model:

$$U_{itj} = V(\mathbf{x_{itj}}, \mathbf{z_i}) + \epsilon_{itj} \tag{2}$$

where $\mathbf{x_{itj}}$ is a vector of time-varying attributes of the J choices and HIU. Utility also depends on fixed attributes of the HIU ($\mathbf{z_i}$), and an unobservable component ϵ_{itj}. Recall, as well, that a function $B(\cdot)$ maps utility from choice j to $r_{ij} = P(y_{it} = j)$, the probability of individual i selecting choice j.

The specification of choice probabilities via a link function to an underlying utility maximization model is the theoretical chassis for most major microsimulation models of the U.S. health care system. This includes models used by the Congressional Budget Office (CB), the RAND Corporation, and the Urban Institute, among others.

For example, the CBO model utilizes a similar underlying utility equation:

$$U_{ij} = \beta_1 V_{ij} + \epsilon_{ij} \tag{3}$$

where the parameter β rescales utility into dollar terms.

The CBO health reform model similarly defines a link function $B(\cdot)$ converting utility to choice probabilities based a on nested logit in which individuals first select the *type* of insurance they will have (e.g., employer, non-group, public, or uninsured) and then conditional on that choice, select among plans within that type. One difference, however, is that the CBO only models marginal changes in coverage, not transitions as we do here. In that sense, the CBO model yields estimates analogous to those derived from repeated cross section data over a 10-year budget horizon.

With the underlying utility structure specified, policy changes are then modeled to affect utility/take-up through their impact on prices, quality, offers of employment-based insurance, etc. In a microsimulation model, this affects the systematic component of utility (i.e., V_{ij}), which is modeled directly using calibrated microdata on individuals and simulated employer choices to offer insurance. Attributes of plans and individuals in the microdata are adjusted to reflect the modeled reform scenario. Specific parameters in the systematic component of the CBO microsimulation model are summarized in the equation below.

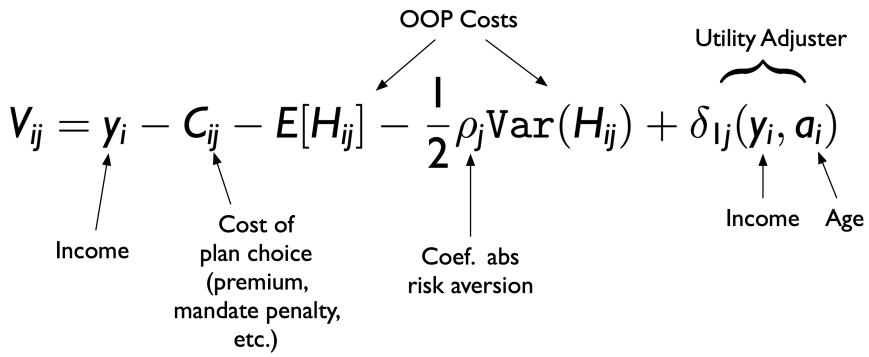

Similarly, in an elasticity-based microsimulation model—which the CBO used prior to 2018—price changes for each of the J insurance options are simulated for units in the microdata. Elasticities and further adjustments (e.g., income effects) are then applied to derive new choice probabilities. Aggregated choice probabilities, along with attributes of individuals (e.g., health status) and policy (e.g., subsidy schedules) are the building blocks for other modeled outcome changes (e.g., cost of subsidies, premiums, etc.). For example, the Gruber Microsimulation Model, which was used by the White House and Congress to model the ACA, used an underlying reduced-form take-up equation with the following form:

$$P(y_{it} = j) = (\texttt{Constant} + \texttt{Elasticity} \times \texttt{Percent Price Change} \times \texttt{Income Effect}) * \texttt{Income Adjustment}$$

1.4 Modeling $\boldsymbol{R(1)}$ and the MVPF: A Reduced-Form Approach

A key feature of the modeling structure developed here is that researchers can simply take the off-the-shelf estimates of $\tilde{\boldsymbol{p}}$ and $\boldsymbol{R(0)}$ from above, and then derive their own estimates of $\boldsymbol{R(1)}$ from reduced form evidence (i.e., without the need for a microsimulation model). Additional parameters summarizing the costs, incidence and valuation of public vs. private plans can then be brought to bear to estimate the comparative welfare effects of different mechanisms for coverage reform.

This section details how we derive estimates of $\boldsymbol{R(1)}$ and the MVPF for two coverage expansion scenarios. We first discuss how we use regression discontinuity evidence on take-up, costs, and individuals' valuation of subsidized private plans to model a hypothetical expansion of \$25/month plans to all low-income individuals under 150% FPL. We then outline an identification strategy for estimating coverage transitions under a public program expansion in nonexpansion states. Our estimates of the MVPF for public expansion also rely on literature-based parameters summarizing uncompensated care costs and valuation of public programs from experimental evidence emerging from the Oregon Health Insurance Experiment.

1.4.1 Subsidies for Private Coverage

Estimates of take-up, valuation and costs of expanding coverage via subsidized private plans are derived from Finkelstein, Hendren, and Shepard (2019). Fundamentally, these estimates are identified off of a regression discontinuity design that keys off of discrete changes coverage take-up and enrollee costs at income boundaries set by the underlying subsidy design adopted by the state of Massachusetts after its 2006 health reforms.

Using replication data provided in Finkelstein, Hendren, and Shepard (2019), we fit the underlying RD regressions and successfully replicated Figure 13 in the final published manuscript. To model a private plan with comparable generosity to Medicaid, we focus on replicated estimates for the "High Generosity" plan.

Figure 6 shows the replicated willingness-to-pay (WTP), marginal cost and average cost curves at various monthly (out-of-pocket) premium levels. By moving from a premium value on the y-axis to the WTP curve and down to the x-axis, we obtain an estimate of the fraction of the market that would take up coverage. For example, in our application we model a \$25/month plan (with the remainder of the total covered by federal subsidies). The lines overlaid on the plot show that a \$25/month subsidy would induce takeup (s^*) among about 85% of the eligible market. This results in costs for the marginal enrollee around \$155/month (black triangle).

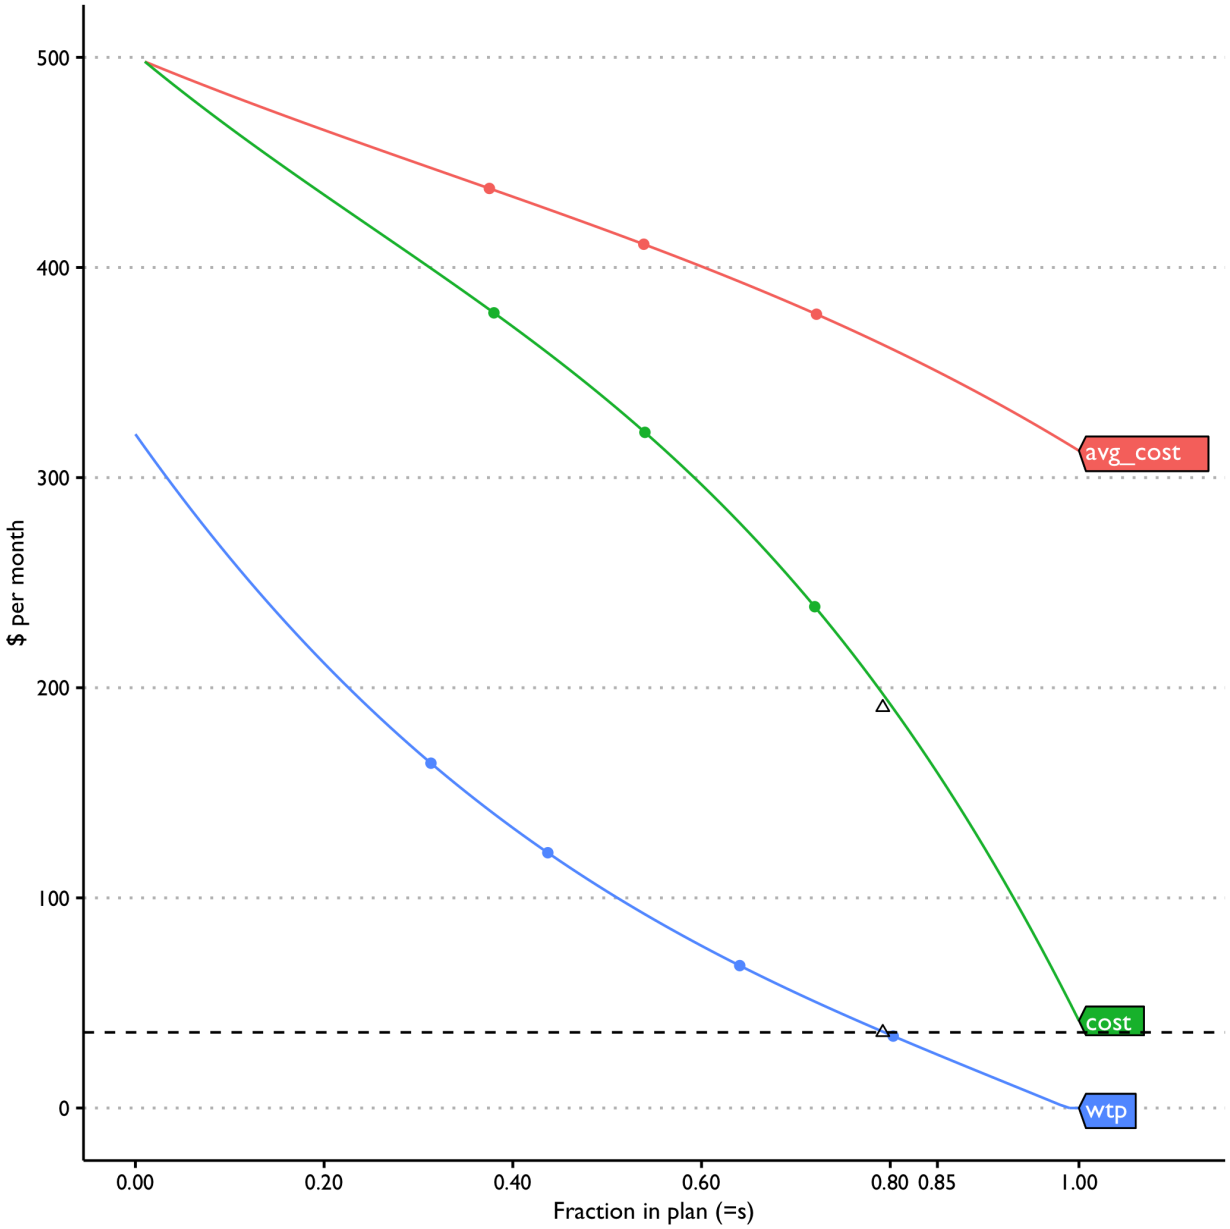

Figure 5: Replication of Figure 13 in Finkelstein, Hendren and Shepard (AER 2020)

For our baseline analyses we use the point-estimate for coverage takeup ($s^* = 0.85$) among those eligible, i.e.,

uninsured individuals <150% FPL. We assume that 70% of the uninsured are eligible, as some may not be U.S. citizens or have legal status that would make them eligible for a federal tax credit, and some may have an offer of insurance through an employer that, under the ACA, would make them ineligible for premium tax credits. The final transition probability matrix adjusts transitions from uninsured-to-non-group and uninsured-to-uninsured to reflect these takeup and eligibility assumptions:

$$\boldsymbol{R}(1)_{subsidy,expansion} = \begin{pmatrix} 0.767 & 0.037 & 0.143 & 0.054 \\ 0.201 & 0.57 & 0.139 & 0.09 \\ 0.102 & 0.019 & 0.838 & 0.041 \\ 0.079 & 0.417 & 0.134 & 0.37 \end{pmatrix}$$

$$\boldsymbol{R}(1)_{subsidy,non-expansion} = \begin{pmatrix} 0.778 & 0.002 & 0.103 & 0.117 \\ 0.27 & 0.53 & 0.154 & 0.046 \\ 0.078 & 0.024 & 0.726 & 0.172 \\ 0.074 & 0.433 & 0.075 & 0.418 \end{pmatrix}$$

In addition to modeling takeup, we also attach state values to uninsured-to-private insurance transitions to model the MVPF of expanding coverage to the uninsured via subsidies to purchase private plans.[3] To do this we use data, inputs and assumptions from the final published version of Finkelstein, Hendren, and Shepard (2019) and carry out the MVPF estimation exercise in the original working paper version (Finkelstein, Hendren, and Shepard 2017).[4] Our modeling of the MVPF includes parameterizations that account for the costs and incidence of uncompensated care on the uninsured vs. on third-parties such as higher-income taxpayers, providers, and the government.

The MVPF for an additional dollar spent to subsidize insurance is modeled using the following equation:

$$MVPF_{subsidy} = \frac{s^* + (1-\gamma)\eta\frac{C_U(s^*)}{-W'_H(s^*)}}{s^* + \frac{C_H(s^*) - \gamma C_U(s^*) - W_H(s^*)}{-W'_H(s^*)}} \tag{4}$$

While details on the derivation of Equation 4 are left to Finkelstein, Hendren, and Shepard (2017) and Finkelstein, Hendren, and Shepard (2019), we will briefly walk through the parameters here. First, s^* is the fraction of the eligible market enrolled. In the absence of uncompensated care, this would represent the marginal benefit of the subsidy: reducing premiums by \$1 is valued at \$1 for all inframarginal enrollees, while the marginal enrolles who choose to insure due to the additional subsidy are indifferent between being insured and uninsured.

The well-documented presence of uncompensated care in the U.S. health care system means that there are additional benefits to third parties from gains in coverage induced by the subsidy. This is captured by the second term in the numerator of Equation 4. Define $C_U(s^*)$ as the cost of uncompensated care:

[3] In principle the entire exercise could be expanded to model additional insurance transitions that could occur in the presence of additional subsidies—for example, transitions from employer-sponsored insurance to subsidized private coverage. This type of modeling was beyond the scope of this study, however, since it would require additional estimation of the welfare costs and benefits of "crowd-out" among private plans.

[4] Correspondance with the authors confirmed that the MVPF exercise was excluded from the final publication due to space constraints and not due to any underlying issues with the MVPF analyses.

$$C_U(s) = (1-x)\frac{C_H(s)}{1+\phi}$$

where x is the share of the uninsureds' health care costs that they pay out of pocket and ϕ is the percentage increase in costs induced via moral hazard effects of becoming insured. Our baseline values assume $x = 0.2$ and $\phi = 0.25$, which are based on parameters derived from the Oregon Health Insurance Experiment (Finkelstein, Hendren, and Luttmer 2019). However, as we describe below, our probabilistic sensitivity exercises explore a range of values for these parameters.

Equation 4 also includes γ, which summarizes the fraction of uncompensated care borne by the government, and η, which is the social marginal utility of income among non-government third-parties that bear the remaining incidence (e.g., hospitals, the privately insured, etc.). In short, η captures the idea that society is indifferent between \$1 in additional subsidy going to low-income uninsured individuals and \$$\eta$ going to third-parties that currently help finance uncompensated care. η values of 1.0 indicate that the remaining incidence of uncompensated care is on the uninsured themselves, while values less than 1 capture the idea that the incidence falls on more affluent indiviudals (e.g., because their social marginal utility of income is less than subsidy recipients, or because there is an efficiency cost to taxing affluent taxpayers to transfer \$1 in subsidy to lower-income uninsured individuals).

Finally, $W_H(s^*)$ is the willingess to pay when s^* of the eligible market is insured, $C_H(s^*)$ is the cost of marginal enrollees, and $\frac{ds^*}{-dp_H} = \frac{1}{-W'_H(s^*)}$ is the change in the market size from a \$1 monthly premium subsidy. These values are derived from points on the curve in Figure 6.

Applying this equation to the data and estimates from Finkelstein, Hendren, and Shepard (2019) yields the following replication of Figure 14 in the original working paper (Finkelstein, Hendren, and Shepard 2017):

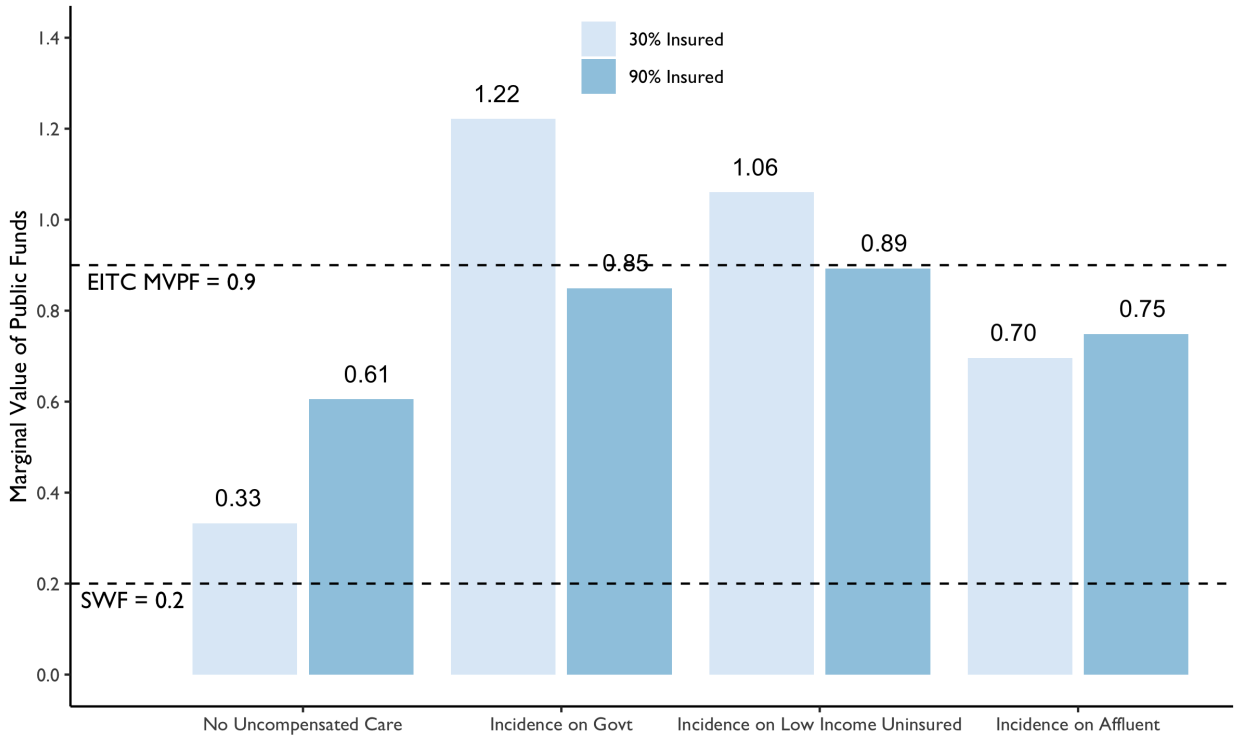

Figure 6: Replication of Figure 14 in Finkelstein, Hendren and Shepard (NBER 2017)

These are just point estimates, however—the estimates above to not provide a sense of how much uncertainty there is as a result of sampling error in the underlying RD regression, or as a result of assumptions made over uncompensated care incidence or other parameters.

To provide a sense of this uncertainty, Figure 7 adds in 95% confidence regions based on simulating multivariate normal draws from the underlying RD regression model. As the figure makes clear, there is less uncertainty around point estimates closer to the observed data points (dark points in figure)—and considerably more at the highest and lowest premium values.

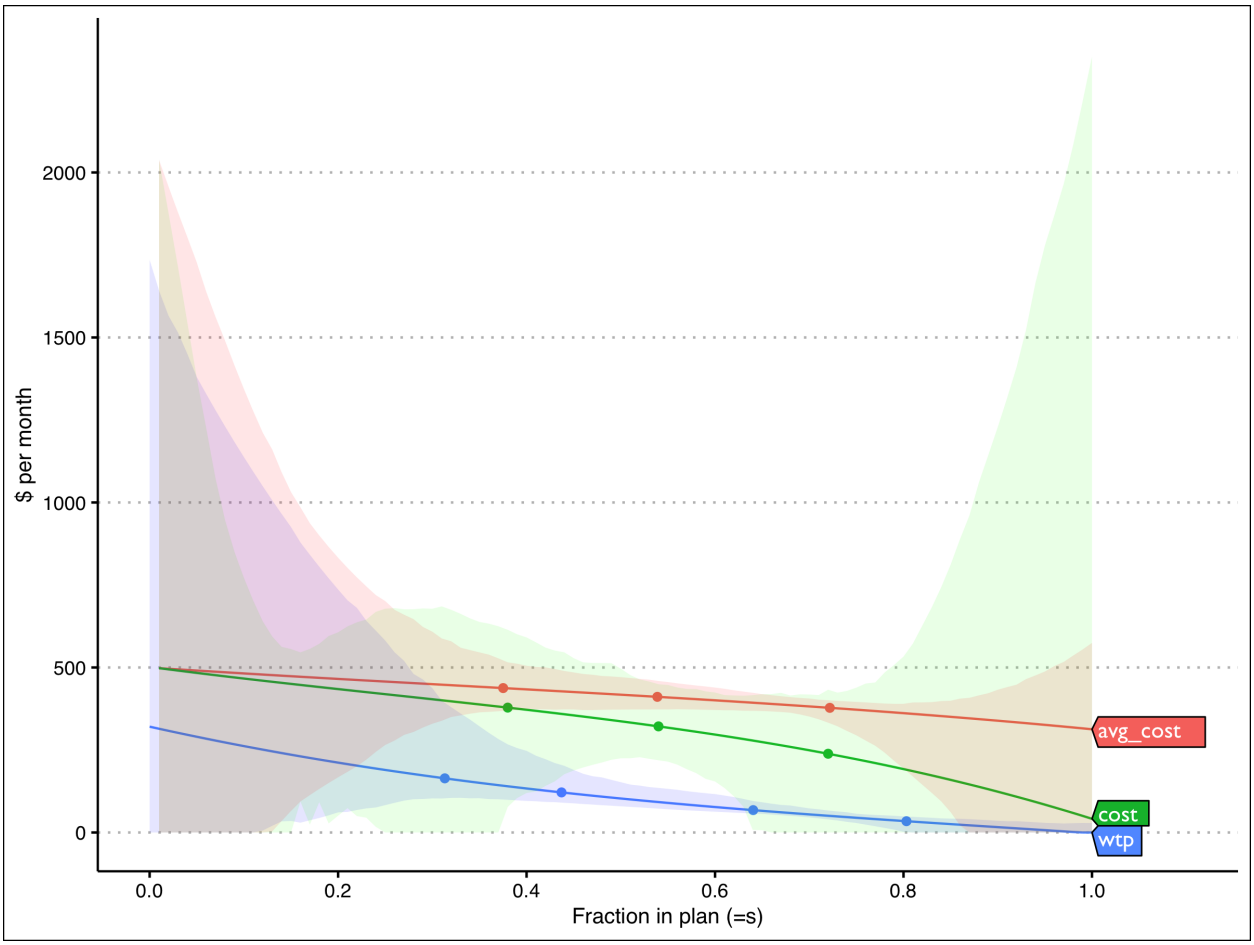

Figure 7: Replication of Figure 12 in Finkelstein, Hendren and Shepard (AER 2019), with 95% confidence regions included

1.4.2 Expansion of Public Coverage

Estimates of changes in transition probabilities as a result of expansion of public coverage are derived via a differences-in-differences design. The design employed here leverages state variation in coverage expansion from 2014-2015, and is based on data from the 2014 SIPP panel. Specific details on the underlying transitions differences-in-differences approach can be found in the manuscript and supplementary material for J. Graves, Fry, et al. (2020) and J. A. Graves et al. (2020).

The basis for identification of changes in transition probabilities stems from standard differences-in-differences assumptions embedded within a similar Markov-based discrete transitions model. For the discrete coverage type outcome, define vectors summarizing baseline occupancy separately by expansion group ($z \in 0, 1$) $\mathbf{p}_{\mathrm{z}}^{\mathrm{pre}} = \left\{ p_{\mathbf{1},\mathrm{z}}^{\mathrm{t=pre}}, \ldots, p_{\mathrm{J},\mathrm{z}}^{\mathrm{t=pre}} \right\}$, where $p_{\mathrm{j},\mathrm{z}}^{\mathrm{t=pre}}$ is the probability of a unit in expansion group z having outcome category j in the pre-expansion period.

We are interested in π_1, the average treatment effect on the treated (ATT) for coverage expansion:

$$\pi_1 = \mathbf{p}_1^{\mathbf{post}}(1) - \mathbf{p}_1^{\mathbf{post}}(0)$$

As in the model discussed in the main text, the post-expansion occupancy vector can be expressed as the product of the pre-period occupancy vector and a transition probability matrix:

$$\mathrm{p}_{\mathrm{z}}^{\mathrm{post}} = \mathrm{p}_{\mathrm{z}}^{\mathrm{pre}\prime}\mathbf{R}_{\mathrm{z}}$$

In the ATT equation above, $\mathbf{p}_1^{\mathbf{post}}(0)$ is unobserved, but we can appeal to the assumptions of differences-in-differences to identify it. More formally, we assume that the untreated (nonexpansion) group's change in outcomes can stand in for the counterfactual in the treated (expansion) group:

$$\mathbf{p}_0^{post}(0) - \mathbf{p}_0^{pre}(0) = \mathbf{p}_1^{post}(0) - \mathbf{p}_1^{pre}(0)$$

Rearranging gives us

$$\mathbf{p}_1^{post}(0) = \mathbf{p}_0^{post}(0) - \mathbf{p}_0^{pre}(0) + \mathbf{p}_1^{pre}(0)$$

Which is equivalent to observed quantities in the pre and post periods for the expansion and nonexpansion groups:

$$\mathbf{p}_1^{post}(0) = \mathbf{p}_0^{post} - \mathbf{p}_0^{pre} + \mathbf{p}_1^{pre}$$

We now can swap in the representation of the post-period occupancy vector using the pre-period vector multiplied by the transition probability matrix:

$$\mathbf{p}_1^{post}(0) = \mathbf{p}_0^{pre\prime}\mathbf{R}_0 - \mathbf{p}_0^{pre} + \mathbf{p}_1^{pre}$$

Swapping this into the ATT equation yields:

$$\pi_1 = \mathbf{p}_1^{pre\prime}\mathbf{R}_1 - (\mathbf{p}_0^{pre\prime}\mathbf{R}_0 - \mathbf{p}_0^{pre} + \mathbf{p}_1^{pre}) \tag{5}$$

In this representation of the ATT, estimates of $\mathbf{p}_1^{pre}, \mathbf{p}_0^{pre}, \mathbf{R}_1$, and $\mathbf{R}_0$ are all that are needed to recover marginal DD treatment effect estimates for the coverage outcome.

Estimates $\hat{\mathbf{p}}_1^{pre}$ and $\hat{\mathbf{p}}_0^{pre}$ can be tabulated directly, since they are just ex ante occupancy vectors summarizing the fraction of individuals in each category in the pre-expansion period.

To estimate $\hat{\mathbf{R}}_0$ and $\hat{\mathbf{R}}_1$ we will draw on a linear DD regression model fit to repeated measures on insurance coverage as measured in SIPP data. To begin, note that a standard DD regression model identifies marginal outcome changes associated with treatment:

$$h(E(Y_{itj})) = \beta_0 + \beta_{j1}Z_i + \beta_{j2}\texttt{Post}_t + \beta_{j3}(Z_i \times \texttt{Post}_t)$$

Where Z_i is a binary indicator for treatment (i.e., expansion), $\texttt{Post}_t$ is an indicator for whether the individual's observation at time period t is in the post-expanison period, and $h(.)$ is a link function

appropriate for the outcome.

Since we have repeated (longitudinal) measures on the outcome for each individual, we augment this model to fit *transitions* in the outcome by interacting the individual's initial coverage status with the differences-in-differences model matrix variables:

$$h(E(Y_{itj})) = \beta_{j0}+\beta_{j1}Z_i+\beta_{j2}\text{Post}_t+\beta_{j3}(Z_i\times\text{Post}_t)+\boldsymbol{\beta_4'}\mathbf{Y}_{i,pre}+\boldsymbol{\beta_5'}\mathbf{Y}_{i,pre}Z_i+\boldsymbol{\beta_6'}\mathbf{Y}_{i,pre}\text{Post}_t+\boldsymbol{\beta_7'}\mathbf{Y}_{i,pre}(Z_i\times\text{Post}_t) \tag{6}$$

where $\boldsymbol{Y}_{i,pre}$ is a vector of indicators for the coverage category at baseline ($\boldsymbol{Y}_{i,pre} = (Y_{i,1,pre}, \ldots, Y_{i,J,pre})'$).

Having fit Equation 6 to paired (two-period) data on individuals in the 2014 SIPP (and comparing coverage transitions between January 2013 and December 2014), we next use recycled predictions to compute the modeled transition probabilities for all $J \times J$ transition types among expansion and nonexpansion state samples. We arrange these modeled transitions into transition probability matrices, and then plug into Equation 5. This yields the following collection of differences-in-differences estimates of the causal effect of Medicaid expanison on coverage transitions:

$$\boldsymbol{R}^{DD} = \begin{pmatrix} 0.0488 & -0.0428 & 0.0507 & -0.0567 \\ -0.0413 & 0.0222 & 0.1111 & -0.0921 \\ -0.0139 & -0.0078 & -0.0185 & 0.0402 \\ -0.0084 & -0.0386 & 0.2288 & -0.1818 \end{pmatrix}$$

Based on this, we model $\boldsymbol{R}(1)$ for nonexpansion states as $\boldsymbol{R}(1) = \boldsymbol{R}(0) + \boldsymbol{R}^{DD}$. For expansion states we do not add in DD estimates since those states already have expanded Medicaid coverage available to those <138% FPL.

We summarize uncertainty around estimates of $\boldsymbol{R}^{DD}$ using a randomization inference procedure as described in J. A. Graves et al. (2020). In short, this procedure re-estimates the underlying differences-in-differences model using a series of 1,000 permuted draws of state expansion status. The collection of K estimates $\boldsymbol{R}^{DD(k)}$ allows us to calculate p-values and otherwise approximate a sampling distribution of estimates.

Figure 8 summarizes robustness checks based on the transitions differences-in-differences model. Specifically, to construct Figure 8, we draw on monthly insurance coverage data in the 2014 SIPP panel and fit a series of two-period transition models based on Equation 6. Each of these models uses January 2013 as the baseline month, however the "post" period month is then allowed to vary from February 2013 through December 2014. As seen in the figure, "post" periods defined up through the end of 2014 (before the actual expansions occurred) yield DD results that hover around zero. As expected,those defined after the actual Medicaid expansion date in January 2014 show large changes in the probability of reporting public coverage.

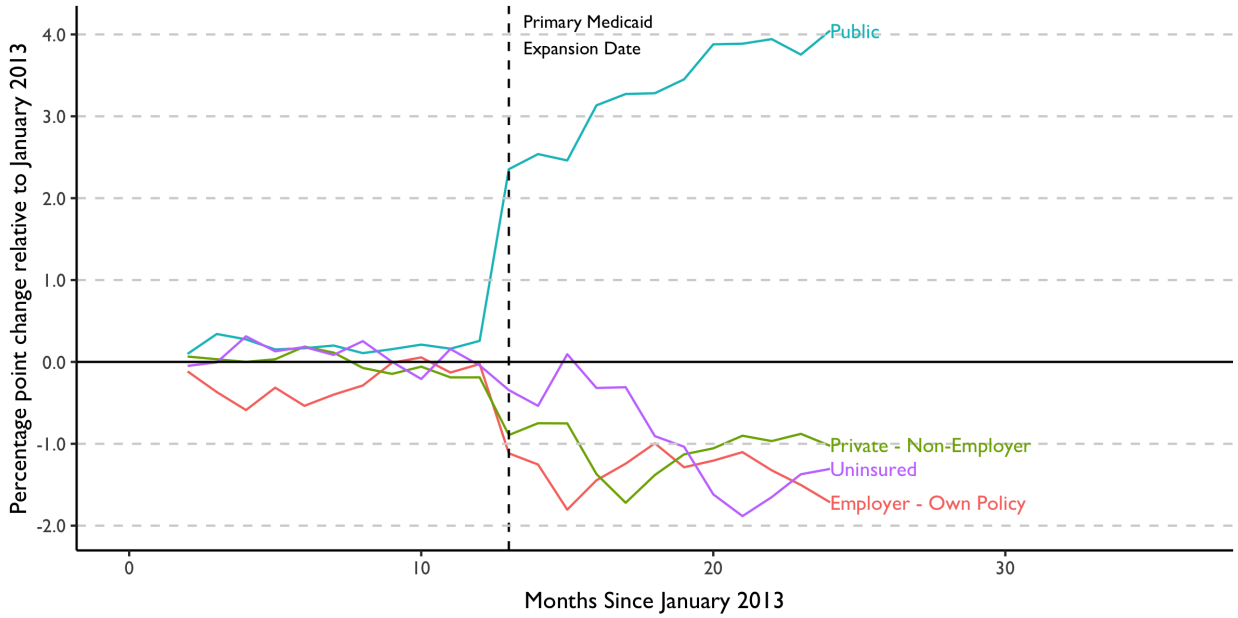

Figure 8: Marginal Coverage Changes Relative to January 2013 for Transitions Differences-in-Differences Model Fit to Alternate Post-Expansion Dates

With estimates of $\boldsymbol{R(1)}_{Medicaid}$ identified and estimated, we next turn to assigning state values to model the MVPF of expanding public coverage. For this exercise we rely on the approach outlined in Finkelstein, Hendren, and Luttmer (2015).

As noted in the main text, the welfare benefits of public coverage expansion are summarized as:

$$W_m = \theta\gamma + (1-\theta)\big(\gamma + \eta\frac{N}{G}\big) \tag{7}$$

where θ is the government incidence of uncompensated care and η is the social marginal utility of income, γ is recipients' willingness to pay for Medicaid (defined below), and N is the monetary transfer of Medicaid to external parties. This monetary transfer is defined as the gross cost of Medicaid G minus its net cost (C) to the government, i.e., $N = G - C$. Net costs (C) are defined as

$$C = (G - G_U) + OOP_U$$

where G_U is government spending on Medicaid recipients if they were uninsured, and OOP_U is out-of-pocket spending among these individuals when uninsured. Intuitively, C captures the observation that net costs to the goverment equal gross Medicaid costs minus what the government would spend on recipients if they were uninsured, plus what those recipients would spend while uninsured, since those costs would be covered under the Medicaid benefit.

Equation 7 includes γ, which quantifies recipients' willingness to pay for Medicaid benefits. Following FHL, γ is modeled as

$$\gamma = \overbrace{\underbrace{\left[\left(\frac{OOP_U}{G_U}\right) - \left(\frac{OOP_M}{G}\right)\right]}_{\text{Change in OOP price for medical care}} \cdot \underbrace{\frac{1}{2}(G_U + G)}_{\text{Linear approximation}}}^{\text{Transfer Term}} + \overbrace{I}^{\text{Pure insurance term}} \tag{8}$$

where OOP_U and OOP_M are recipients' out of pocket costs when uninsured and insured under Medicaid, respectively. Similarly, G_U and G are government spending on the uninsured and on Medicaid recipients. The pure insurance term, I, in Equation 8 enters the application as a scalar value based on the range of estimates in FHL.[5]

Finally, the welfare costs of in-kind benefits are represented by

$$C_m = \theta C + (1 - \theta)G \tag{9}$$

where C is the net cost of Medicaid and G is the gross cost of Medicaid, as defined above. In this equation, the first term captures the observation that if the incidence of uncompensated care falls on the government, costs are captured by the net costs after accounting for existing government financing for uncompensated care. By comparison, the second term captures costs when the incidence falls on external (non-government) parties. In that case, costs to the government are the full gross costs.

As shown in Figure 9, the above equations successfully replicate—to within a small amount of error—the main welfare results in Figure 14 of Finkelstein, Hendren and Shepard (2019).

[5] A nice feature of this study's approach is that we can parameterize the model based on a minimal number of parameters. Then, using the VOI methods employed here, we can identify the paramters with the highest amount of leverage. To the extent a parameter such as I proves influential, we can then re-parametrize the model by peeling back the layers of the I onion. In this way we can use the VOI exercise as both a guide for future research on uncertain parameters, and for model building itself.

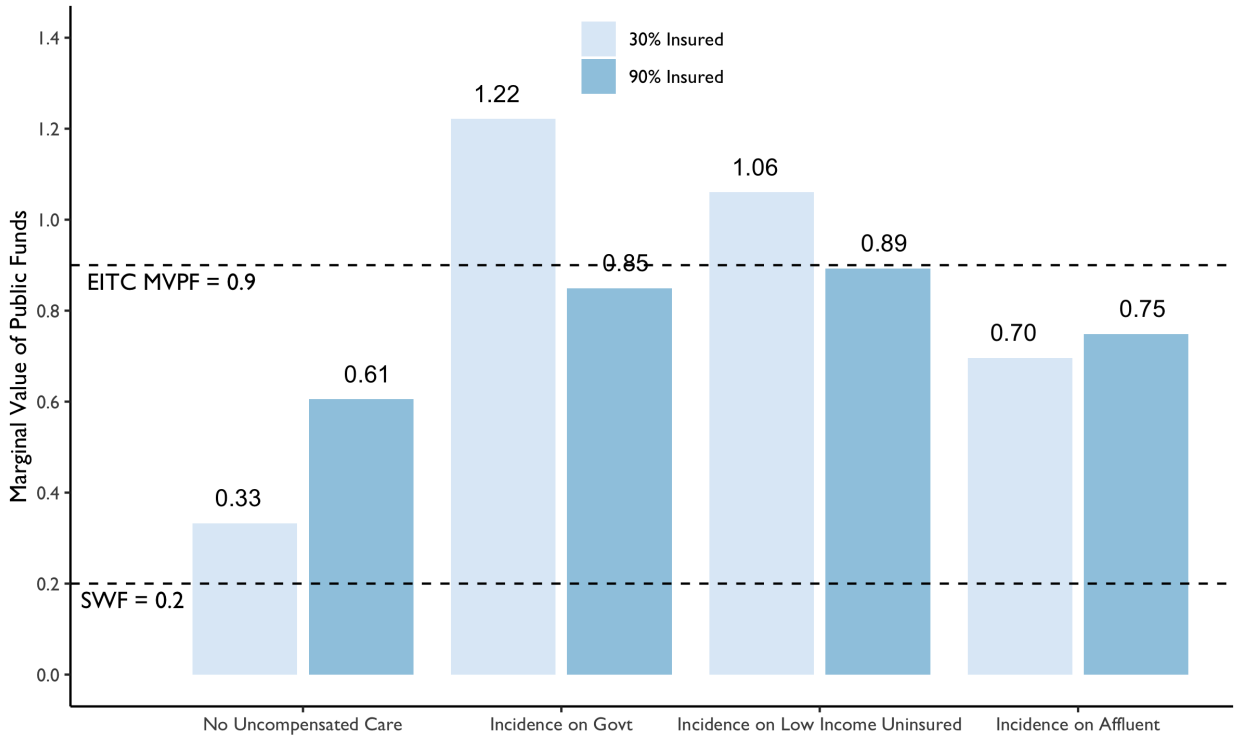

Figure 9: Replication of Main Results in Finkelstein, Hendren and Shepard (2019)

References

Asgharian, Masoud, Cyr Emile M'Lan, and David B. Wolfson. 2002. "Length-Biased Sampling with Right Censoring: An Unconditional Approach." *Journal of the American Statistical Association* 97 (457): 201–9.

de Uña-Alvarez, Jacobo, M. Soledad Otero-Giráldez, and Gema Alvarez-Llorente. 2003. "Estimation Under Length-Bias and Right-Censoring: An Application to Unemployment Duration Analysis for Married Women." *Journal of Applied Statistics* 30 (3): 283–91.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6): 2836–74. https://doi.org/10.1086/702238.

Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." National Bureau of Economic Research.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2017. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." National Bureau of Economic Research.

———. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530–67.

Graves, John A., Laura A. Hatfield, William Blot, Nancy L. Keating, and J. Michael McWilliams. 2020. "Medicaid Expansion Slowed Rates of Health Decline for Low-Income Adults in Southern States: An Analysis

of the Impact of Medicaid Expansion on the Self-Reported Health of Low-Income Older Nonelderly Adults Living in the South." *Health Affairs* 39 (1): 67–76.

Graves, John A., and Sayeh S. Nikpay. 2017. "The Changing Dynamics of US Health Insurance and Implications for the Future of the Affordable Care Act." *Health Affairs* 36 (2): 297–305. https://doi.org/10.1377/hlthaff.2016.1165.

Graves, John, Carrie Fry, Michael McWilliams, and Laura A Hatfield. 2020. "Difference-in-Difference Estimation for Transitions in Discrete Outcomes: Insurance Transitions After the Affordable Care Act's Medicaid Expansion." Working Paper.

Graves, John, Shawn Garbett, Zilu Zhou, Jonathan Schildcrout, and Josh Peterson. 2020. "Comparison of Decision Modeling Approaches for Health Technology and Policy Evaluation | medRxiv." https://www.medrxiv.org/content/10.1101/2020.05.26.20113845v1.

Raftery, Adrian E., and Le Bao. 2010. "Estimating and Projecting Trends in HIV/AIDS Generalized Epidemics Using Incremental Mixture Importance Sampling." *Biometrics* 66 (4): 1162–73.

Rutter, Carolyn M., Jonathan Ozik, Maria DeYoreo, and Nicholson Collier. 2019. "Microsimulation Model Calibration Using Incremental Mixture Approximate Bayesian Computation." *The Annals of Applied Statistics* 13 (4): 2189–2212.