American Economic Review 2019, 109(4): 1461–1485
https://doi.org/10.1257/aer.20180325

Targeting with In-Kind Transfers: Evidence from Medicaid Home Care[†]

By Ethan M. J. Lieber and Lee M. Lockwood*

Making a transfer in kind reduces its value to recipients but can improve targeting. We develop an approach to quantifying this trade-off and apply it to home care. Using randomized experiments by Medicaid, we find that in-kind provision significantly reduces the value of the transfer to recipients while targeting a small fraction of the eligible population that is sicker and has fewer informal caregivers than the average eligible. Under a wide range of assumptions within a standard model, the targeting benefit exceeds the distortion cost. This highlights an important cost of recent reforms toward more flexible benefits. (*JEL* D82, H51, H75, I18, I38)

In-kind transfers are a ubiquitous feature of government programs, private contracts, and charitable giving. In the United States, government spending on in-kind programs exceeds 12 percent of GDP and spending on in-kind health programs alone exceeds $1 trillion per year (Currie and Gahvari 2008, Centers for Medicare and Medicaid Services 2017). In domestic policy, foreign aid, and charitable giving, there are active debates about the desirability of flexible benefits such as direct cash transfers and universal basic income programs versus restrictive in-kind transfers of food, housing, medical care, and other goods.

Central to these debates is a trade-off inherent to in-kind transfers. In-kind provision has a fundamental cost: recipients would prefer an equal-cost cash transfer. But this cost is linked to an important potential benefit. In-kind provision can better target desired recipients by leading some people to take up more benefits than others (Nichols and Zeckhauser 1982). In the context of insurance, if someone values a particular good more in states of the world in which marginal utility is higher, an

*Lieber: Department of Economics, University of Notre Dame, 3049 Nanovic Hall, Notre Dame, IN 46556 (email: Ethan.Lieber.2@nd.edu); Lockwood: Department of Economics, University of Virginia, 248 McCormick Road, Charlottesville, VA 22903, and NBER (email: leelockwood@virginia.edu). Liran Einav was coeditor for this article. We are grateful to the anonymous referees, Norma Coe, Manasi Deshpande, Amy Finkelstein, Gopi Shah Goda, Tami Gurley-Calvez, Seema Jayachandran, Brian Melzer, Matt Notowidigdo, Mike Powell, Emmanuel Saez, Diane Schanzenbach, Jesse Shapiro, Courtney Van Houtven, our department colleagues, and many seminar participants for helpful comments. We thank José Carreno, Vishal Kamat, Max Rong, and Zeyu Wang for excellent research assistance. The research reported herein was performed pursuant to a grant from the US Social Security Administration (SSA) through the Michigan Retirement Research Center (Grant 5 RRC08098401-05), funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of the Social Security Administration, any agency of the Federal Government, or the Michigan Retirement Research Center. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to https://doi.org/10.1257/aer.20180325 to visit the article page for additional materials and author disclosure statement(s).

in-kind transfer of that good can help concentrate benefits in those states and thereby better insure the risk. Although these costs and benefits of in-kind provision are crucial determinants of optimal policy, little is known about their relative magnitudes across a wide range of important contexts.

In this paper, we develop an approach to quantifying this core trade-off of in-kind provision and apply it to home care. Home care helps people who have chronic health problems with tasks such as eating, dressing, and bathing. Its value, including care from family and friends ("informal care") as well as professional caregivers ("formal care"), is thought to exceed $200 billion per year (Arno, Levine, and Memmott 1999). Traditionally, home care benefits have been provided as in-kind formal care. But following Medicaid's large-scale Cash and Counseling experiments in the late 1990s, many states reformed their home care programs to make benefits more flexible and cash-like (National Conference of State Legislatures 2007).

Our approach to quantifying the welfare effect of in-kind provision involves three main ingredients. The first is the moral hazard effect, the extent to which in-kind provision increases consumption of the good. The greater this increase, the lower the value to recipients of the in-kind benefit relative to its cost. Using the randomized assignment of in-kind versus near-cash benefits in the Cash and Counseling experiments, we estimate that in-kind provision increases formal care consumption among those consuming formal care by 25 hours per week, nearly twice the average consumption in the benefit-eligible population. This suggests that many recipients value the in-kind benefit far below its cost. Our estimates imply that a recipient of the average in-kind transfer in the experiment values it at 28 percent of its cost.

The second ingredient is the distribution of consumption of the good within benefit-eligible states of the world. The greater the heterogeneity in consumption of the good, the greater the extent to which in-kind provision concentrates transfers. Using nationally representative data on the eligible population, we find considerable heterogeneity in consumption of formal care. While 63 percent of those eligible do not consume any formal care, among those who do, there is a long right tail. An individual at the ninety-fifth percentile receives around-the-clock care, which at the average hourly price of $15 amounts to about $131,000 per year (Genworth Financial 2005).

The third ingredient is the link between consumption of the good and the marginal utility of income. The stronger this link, the more valuable it is to shift resources to the states of the world targeted by in-kind provision. In our context, this link is likely strong. Greater costs of coping with bad health leave fewer resources for non-health consumption, which tends to increase marginal utility. Empirically, we find that in-kind provision sharply concentrates transfers on a small fraction of the eligible population that has a greater demand for formal care, is sicker, and has fewer informal caregivers than the average eligible. To the extent that such recipients tend to have relatively high marginal utility, in-kind provision could significantly improve insurance.

These results suggest that designers of home care benefits face a stark trade-off. Restrictive in-kind benefits are valued far less than their cost ex post, but they sharply concentrate transfers in what appear to be relatively high-marginal utility states. We combine our reduced-form estimates with a structural model to quantify this trade-off in a stylized expected utility framework. Under a wide range of

assumptions, the optimal contract involves a large in-kind component and delivers substantial welfare gains over a cash-benefit contract.

Our paper complements and extends the literature on barriers to private, voluntary long-term care insurance (see Brown and Finkelstein 2011 for a review). Our findings reveal the critical importance of two factors in determining the welfare effect of any long-term care insurance, whether public or private, voluntary or mandatory: risk within unhealthy states of the world and moral hazard. Although in-kind provision has a large moral hazard cost, the gain from insuring the considerable risk within unhealthy states appears to be even larger. This raises concerns about recent reforms toward cash-like benefits.

Our approach helps link the theoretical and empirical literatures on in-kind transfers, which have been largely disconnected so far (Currie and Gahvari 2008).[1] Methodologically, the equivalence of the effects of an in-kind transfer and a corresponding price subsidy on a recipient's choice set allows us to use ideas from the literatures on optimal taxation and health insurance to quantify a core trade-off of in-kind provision. Substantively, the key feature of in-kind transfers that gives rise to their targeting and distortion effects is that they reduce the recipient's cost of consuming the good over some range of quantities, thereby "loosening" the budget constraint more for recipients who consume more of the good.[2] This feature is shared by a wide range of other policies, including vouchers, conditional cash transfers, benefit programs with ordeals, insurance policies, and commodity taxes and subsidies.

This paper also contributes to the literature on targeting in benefit programs such as housing assistance (Reeder 1985), Medicaid (Cutler and Gruber 1996), Supplemental Security Income (Benitez-Silva, Buchinsky, and Rust 2004), disability insurance (Low and Pistaferri 2015, Deshpande and Li 2017), and food stamps (Finkelstein and Notowidigdo 2018).[3] In many programs, only a small fraction of the eligible population takes up benefits. While low take-up can be undesirable, our findings suggest that it can also significantly increase welfare through better targeting.

I. Approach

This section describes an approach to quantifying the targeting-distortion trade-off of in-kind provision. This trade-off has previously been analyzed in theory (e.g., Nichols and Zeckhauser 1982, Blackorby and Donaldson 1988), using models that

[1] The theoretical branch has investigated potential advantages of in-kind transfers in terms of paternalism (Musgrave 1959), targeting (Nichols and Zeckhauser 1982, Blackorby and Donaldson 1988), tax system efficiency (Munro 1992), and the Samaritan's Dilemma (Bruce and Waldman 1991). Much of the empirical branch estimates the effects of in-kind transfers on consumption (e.g., see Hoynes and Whitmore Schanzenbach 2016 for a review of the effects of food transfers). Other work examines the effects of in-kind transfers on poverty (Smeeding 1977), targeting (Reeder 1985, Cutler and Gruber 1996, Jacoby 1997), and prices (Cunha, De Giorgi, and Jayachandran 2011). Our approach complements those of Finkelstein, Hendren, and Luttmer (2015), who analyze the welfare effect of Medicaid health insurance coverage for prime-age adults.

[2] In-kind transfers that are inframarginal for most potential recipients are unlikely to have large targeting and distortion effects. This seems likely to be the case for the food transfer program in the United States, though there is an ongoing debate about the effect of the transfer on patterns of spending (Hoynes and Whitmore Schanzenbach 2016).

[3] See Currie (2006) for a review. A related literature in the developing world investigates the targeting effects of ordeals (Alatas et al. 2016), subsidized prices (Cohen and Dupas 2010), and delegating authority over the distribution of benefits to local leaders (Alatas et al. 2012, Basurto, Dupas, and Robinson 2017). Kleven and Kopczuk (2011) analyze the role of program complexity in determining take-up.

provide clear insights about the economic factors involved but that are not well suited to empirical implementation. Our approach to empirical implementation has close parallels in the literatures on optimal taxation and health insurance (e.g., Zeckhauser 1970, Mirrlees 1971, Manning and Marquis 1996, Saez 2001). These parallels arise from an economic equivalence: for any in-kind transfer, there is a subsidy that has the same effect on the recipient's budget set.[4]

The key feature of an in-kind transfer that gives rise to the targeting-distortion trade-off is that it reduces the recipient's cost of consuming the transferred good. One consequence is that recipients over-consume the good and value the transfer less than its cost to the provider. This is the moral hazard cost of in-kind provision. Another consequence is that, because the cost reduction is more valuable to someone consuming more of the good, it targets states of the world or types of people with relatively high consumption of the good. If these states or types have relatively high marginal utility, that is a targeting benefit of in-kind provision.

A. *Theory*

An individual faces a risk that potentially affects prices, income, and preferences. The eventual state of the world is uncertain ex ante and non-contractible ex post. As a result, an insurance contract cannot target high-marginal utility states directly by offering larger benefits in those states. Instead, any targeting must be indirect, relying on differential take-up of a single benefit. A natural candidate is an in-kind transfer of an "indicator good," a good consumed in greater quantities in higher-marginal utility states of the world (Nichols and Zeckhauser 1982).

Consider an in-kind transfer with no quantity limit, a linear subsidy. An increase in the subsidy rate reduces the after-subsidy price to consumers. We focus on this case for simplicity and because it best matches our empirical application. Online Appendix Section A.1 analyzes the case with a binding benefit limit; the core trade-off is the same.

Ex post indirect utility in the realized state of the world is

$$v(p,m) \equiv \max_{x_k, x_{-k}} u(x_k, x_{-k}) \quad \text{subject to} \quad p_k x_k + \sum_{i \neq k} p_i x_i \leq m,$$

where p is the vector of prices, m is income, x_k is the good being transferred in kind, and x_{-k} is the vector of all other goods. By the envelope theorem (Roy's Identity), the ex post marginal value of a reduction in the price of good k is the individual's consumption of good k in that state: $\frac{-\partial v(p,m)}{\partial p_k} \Big/ \frac{\partial v(p,m)}{\partial m} = \frac{\lambda x_k}{\lambda} = x_k$, where λ is the marginal utility of income. The individual's consumption of x_k is the amount by which the price reduction "loosens" the individual's budget constraint in that ex post state.

[4] For example, an in-kind transfer that offers recipients up to a fixed amount of the good free of charge has the same effect as a nonlinear subsidy of 100 percent up to that fixed amount and 0 percent thereafter.

The ex ante expected marginal benefit of a reduction in the price of good k is

$$MB = \frac{-\partial E\big(v(p,m)\big)/\partial p_k}{\partial E\big(v(p,m)\big)/\partial m} = \frac{E(\lambda x_k)}{E(\lambda)} = E(x_k) + \operatorname{cov}\big(\hat{\lambda}, x_k\big),$$

where $E\big(v(p,m)\big)$ is expected (indirect) utility and $\hat{\lambda}$ is the marginal utility of income normalized so that its mean is 1.[5] The ex ante value of the price reduction is its mean ex post value, $E(x_k)$, plus a "correction term" for the relationship, if any, between marginal utility and consumption of good k, $\operatorname{cov}\big(\hat{\lambda}, x_k\big)$. This term arises because the subsidy has a larger impact in states with greater consumption of good k. If $\operatorname{cov}\big(\hat{\lambda}, x_k\big) > 0$, the targeting of benefits to states with greater consumption of good k also tends to target benefits to high-marginal utility states, providing insurance that a cash transfer does not. In this case, the ex ante expected marginal benefit of the subsidy exceeds its mean ex post value. This covariance term, the insurance value of the subsidy's differential targeting of states with greater consumption of good k, is the targeting benefit of in-kind provision.

The expected cost to the insurer of the in-kind benefit is $(p_k^0 - p_k)\,E(x_k)$, where p_k^0 is the unsubsidized price and p_k is the consumer's net-of-subsidy price.[6] The marginal cost to the insurer of a reduction in the price of good k is

$$MC = E(x_k) + (p_k^0 - p_k)\,E\left(-\frac{dx_k}{dp_k}\right).$$

The first term is the insurer's additional spending due to the increase in the subsidy rate, holding fixed consumption ("mechanical effect"). The second term is the insurer's additional spending on the subsidy due to the induced change in consumption ("moral hazard effect").[7]

Figure 1 plots the marginal cost, marginal benefit, and mechanical effect of reductions in the price of good k as functions of the subsidy rate, s, where $p_k = (1-s)\,p_k^0$. The targeting benefit is the vertical distance from the mechanical effect to the marginal benefit. The distortion cost is the vertical distance between the mechanical effect and the marginal cost. The marginal cost of the subsidy exceeds its mean ex post value (the mechanical effect) due to moral hazard; in each state, the subsidy is less valuable than an equal-cost cash benefit due to the change in

[5] The second equality comes from the envelope theorem. The final equality comes from noting that $E(\lambda x_k) = E(\lambda)\,E(x_k) + \operatorname{cov}(\lambda, x_k)$ and that $\operatorname{cov}(\lambda, x_k)/E(\lambda) = \operatorname{cov}(\lambda/E(\lambda), x_k) \equiv \operatorname{cov}\big(\hat{\lambda}, x_k\big)$.

[6] This assumes that the supply of every good is perfectly elastic. In this case, an increase in the subsidy reduces the individual's after-subsidy price of good k one-for-one and has no effect on the prices of other goods. This marginal cost does not include second-best considerations from other distortions in the economy, such as substitution from subsidized nursing home care. The problem can therefore be viewed as that of a private insurer offering a stand-alone home care benefit, which would not account for such effects. We discuss the likely impact of some of the main second-best considerations in the conclusion.

[7] The policy counterfactual, in particular who pays for the subsidy, affects the size of the moral hazard effect. The moral hazard effect is the *total* derivative, i.e., the combined effect of the price reduction and any accompanying change in nominal income, $\frac{dx_k}{dp_k} = \frac{\partial x_k(p,m)}{\partial p_k} + \frac{\partial x_k(p,m)}{\partial m} m'(p_k)$, where $x_k(p,m)$ is Marshallian demand for good k and $m'(p_k)$ is the accompanying change in nominal income. We focus on cost-neutral shifts in a mixed in-kind/cash benefit, which pair an increase in the subsidy with a reduction in the uniform cash benefit that holds fixed total spending on recipients. Such cost-neutral shifts isolate the welfare effect of in-kind provision from that of redistribution between recipients and other parties. This means that the marginal cost and marginal benefit are in the same units: income in the hands of recipients.

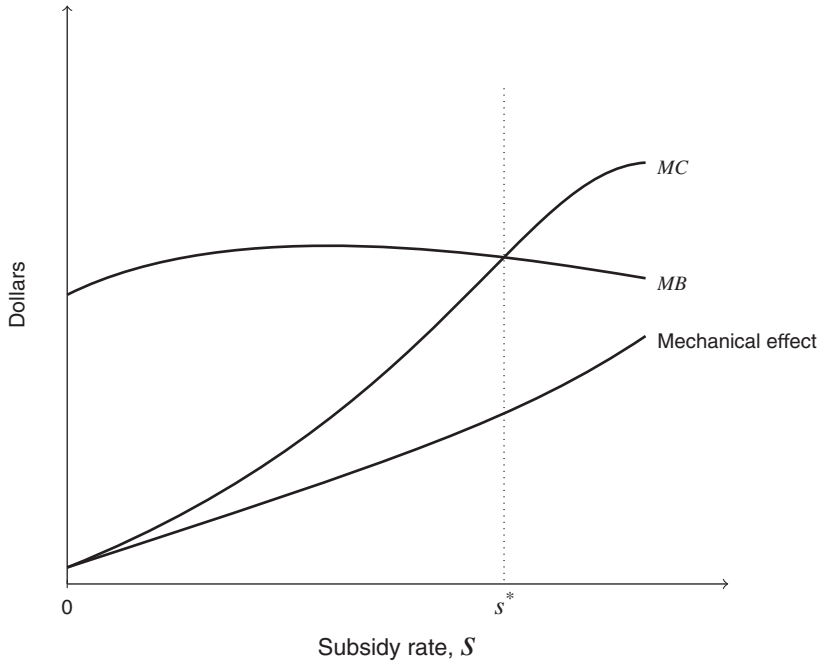

FIGURE 1. COSTS AND BENEFITS OF A SUBSIDY

Notes: *MB* is the marginal benefit of an increase in the subsidy on good k. *MC* is the marginal cost. The *Mechanical effect* ($E(x_k)$) is both (i) what the marginal benefit would be if the targeting benefit were 0 (i.e., if $\text{cov}\left(\hat{\lambda}, x_k\right) = 0$) and (ii) what the marginal cost would be if the moral hazard cost were 0 (i.e., if $E\left(-dx_k/dp_k\right) = 0$).

consumption it induces. At the optimum, the marginal targeting benefit equals the marginal distortion cost and both exceed the mechanical effect. The optimal contract leaves some risk uninsured since the benefit of insuring it is smaller than the cost.

B. *Empirical Implementation*

Our approach to quantifying the targeting-distortion trade-off is based on three ingredients: the price-sensitivity of demand for the good, dx_k/dp_k; the distribution of consumption of the good within benefit-eligible states of the world, $F(x_k)$; and the link between consumption of the good and marginal utility.

The price-sensitivity of demand determines the moral hazard cost of in-kind provision, the excess of the cost of the benefit over its value to recipients ex post.

The remaining two ingredients, the distribution of consumption of the good and the link between consumption of the good and marginal utility, determine the targeting benefit of in-kind provision. Letting σ_X be the standard deviation of variable X across states of the world, we can decompose the marginal targeting benefit as $\text{cov}\left(\hat{\lambda}, x_k\right) = \sigma_{x_k}\sigma_{\hat{\lambda}}\text{corr}\left(\hat{\lambda}, x_k\right)$. The distribution of consumption of the good determines σ_{x_k}. This in turn determines the extent to which in-kind provision concentrates benefits in some states and not others since the ex post marginal benefit of a shift toward in-kind provision is proportional to x_k. Given the distribution of consumption

of the good, the link between consumption of the good and marginal utility determines both the extent of risk ($\sigma_{\hat{\lambda}}$) and the extent to which the states targeted by in-kind provision have relatively high-marginal utility ($\text{corr}(\hat{\lambda}, x_k)$). This decomposition splits the targeting benefit into two parts: (i) the targeting effect, the estimable effect on the distribution of transfers, and (ii) the value of this targeting, which depends on the unobservable link between consumption of the good and marginal utility.

This decomposition isolates assumptions about marginal utility from the rest of the analysis. It facilitates analyses that incrementally build from reduced-form estimations that shed light on the key magnitudes to sufficient-statistics and structural approaches that quantify the net welfare effect. Without any assumptions about marginal utility, straightforward estimations of the price sensitivity of demand and the distribution of consumption reveal the moral hazard cost and targeting effect of in-kind provision. With a qualitative sense of the link between consumption of the good and marginal utility, as presumably exists for a good being transferred in kind, the distribution of consumption is also informative about the extent of risk and the potential targeting benefit of in-kind provision. With a model of marginal utility, the net welfare effect can be quantified as well. The theoretical considerations and empirical evidence that can help inform this important modeling choice will vary by context. Online Appendix Section A.2 discusses the applicability of the approach.

II. Home Care, Medicaid, and the Cash and Counseling Experiments

Chronic health problems are one of the most important risks people face over the life cycle. Roughly 15 percent of Americans over age 50 have at least one person helping them perform activities of daily living (ADL) such as bathing, eating, and dressing (Barczyk and Kredler 2018). Eighty-seven percent of those receiving help live in the community and 74 percent of all care hours occur in private homes (Barczyk and Kredler 2018). Spending on formal home care was $88 billion in 2015, and the total cost of home-based care, including (hard-to-measure) informal care from family and friends, is thought to exceed the total cost of formal long-term care (Arno, Levine, and Memmott 1999, Centers for Medicare and Medicaid Services 2017). Despite the magnitude of this risk, just 10 percent of people 65 and older own private long-term care insurance. As a result, a large share of the costs of long-term care in general and home care in particular are paid by the means-tested Medicaid program.

Medicaid home care programs are an important source of care for many people. In 2013, Medicaid spent $57 billion on the home-based care of more than 3 million recipients. This is about one-half of Medicaid's total spending on long-term care and about two-thirds of all spending on formal home care. Eligibility for Medicaid home care is determined by financial- and health-related criteria. An individual must have sufficiently low income and assets and must have at least two ADL limitations that are expected to last at least 90 days. The traditional Medicaid home care benefit is an in-kind benefit of formal home care from a Medicaid-approved agency. The amount of care an individual can receive free of charge is determined by a "care plan" created by her physician or nurse following a medical examination, though in the specific cases we analyze there does not appear to be a binding upper limit. Online

Appendix Section B discusses evidence on this and provides additional information about Medicaid home care.

In recognition of the importance of informal care and other ways of dealing with chronic health problems, many state Medicaid programs have implemented reforms toward more flexible, cash-like benefits (Doty, Mahoney, and Sciegaj 2010).[8] These programs typically allow recipients to spend their benefits on a wide range of personal care goods and services including assistive devices, home modifications, and, most important, informal care from family or friends. More flexible, cash-like benefits are increasingly common in other countries as well. Germany, France, Italy, Austria, Sweden, and the Netherlands all have long-term care programs that either pay benefits in cash or allow recipients to choose between cash and in-kind benefits (Da Roit and Le Bihan 2010).

An important milestone in the debate about more- versus less-flexible benefits, and an important source of evidence in our paper, is the Cash and Counseling experiments. These were large-scale experiments run by Medicaid programs in Arkansas, Florida, and New Jersey that began in 1998. Participants were drawn primarily from the population of Medicaid home care recipients and were randomized to either the traditional in-kind home care benefit or a near-cash benefit, each with 50 percent probability. Participants randomized to the near-cash benefit could revert to the standard in-kind benefit at any time; those randomized to the in-kind benefit could not switch to the near-cash benefit. Each recipient of the near-cash benefit received a budget for spending on care-related goods and services roughly equal to the cost of the care in her care plan. She also received "counseling" services to help manage her benefit. These services included help with planning how to spend the benefit, hiring and paying caregivers (and paying payroll taxes), and maintaining records. The aim was to make it as easy to receive care for the near-cash group as for the in-kind group. The restriction that the near-cash benefit had to be spent on care-related goods and services was unlikely to be binding for most recipients because of the broad definition of care-related goods and services, especially the inclusion of informal care.[9]

The main goal of the experiments was to test whether recipients could effectively manage their near-cash benefits and receive "enough" care. The results were almost uniformly positive. Members of the near-cash treatment group reported greater satisfaction with their care and with their lives as a whole (Foster et al. 2003, Brown et al. 2007). They also had similar or better health outcomes across a wide range of measures such as mortality, nursing home entry, falls, urinary tract infections, and respiratory infections (Lepidus Carlson et al. 2007). In the official final report on the experiments, Brown et al. (2007) conclude that the near-cash benefit had overwhelmingly positive effects on recipients.

[8] Early versions of the Affordable Care Act included a long-term care insurance program that would have paid cash benefits. This program, the CLASS (Community Living Assistance Services and Supports) Act, was eventually dropped due to concerns about its cost.

[9] The vast majority of participants had been receiving enough informal care at baseline to more than exhaust their benefit. At follow up, 86 percent of recipients of the near-cash benefit used it to pay for informal care (Brown et al. 2007). Online Appendix Section B.2 contains more information about the Cash and Counseling experiments. Online Appendix Table E.1 reports summary statistics of Cash and Counseling participants and balance tests; these provide evidence of a valid randomization. Online Appendix Table E.2 compares Cash and Counseling participants both to the broader population of people eligible for home care benefits and to those who take up Medicaid home care. Our analysis uses data on the 2,470 participants age 65 or older.

III. Moral Hazard Cost of In-Kind Provision

In this section, we estimate the price sensitivity of demand for formal care. This, the first ingredient in our approach, is the key parameter for quantifying the moral hazard cost of in-kind provision. We use the Cash and Counseling experiments, which have two major advantages for this purpose.[10] First, the randomization solves an especially difficult simultaneity problem: many factors that shift the supply of formal care are also likely to shift the demand for formal care by changing the opportunity cost of informal care.[11] Second, the variation in the price of formal care spans the full range most relevant for policy, from zero to the market price.

The experimental results suggest that in-kind provision of home care has a large moral hazard cost. Table 1 shows that being randomized to in-kind benefits doubles average formal care consumption from 7 to 14 hours per week. Figure 2 shows the distributions of formal care consumption for those randomized to the in-kind versus near-cash benefits. In-kind provision increases formal care consumption throughout the distribution, more than doubling both the fraction of people who consume formal care (from 24 to 55 percent) and the fraction who consume more than 20 hours per week (from 9 to 22 percent).

We estimate the price sensitivity of demand for formal care taking into account censoring at zero and imperfect compliance. We account for censoring by treating an individual's observed hours of care, q_i, as the outcome of a censored, latent demand for care, $q_i = \max\{0, q_i^*\}$. We account for imperfect compliance (some people assigned to the near-cash benefit reverted to the traditional in-kind benefit and some left Medicaid home care altogether) by using the randomized assignment as an instrument for the price each participant faced. Participants who receive the near-cash benefit or who leave Medicaid home care face the market price in their state. Participants who receive the in-kind benefit face a price of zero.[12] We estimate the system

$$
\begin{aligned}
q_i^* &= \alpha + \beta p_i + X_i\gamma + \varepsilon_i, \\
q_i &= \max\{0, q_i^*\}, \\
p_i &= \mu_0 + \mu_1 Cash_i + X_i\mu_2 + \nu_i,
\end{aligned}
$$

[10] Previous research on the Cash and Counseling experiments has focused on the distinction between *paid* and *unpaid* home care, where paid home care includes care from family and friends as well as from professionals, so long as the recipient pays for it (e.g., Brown et al. 2007). We focus on the distinction between formal care, provided by professionals, and informal care, provided by family and friends, regardless of whether the recipient pays the caregiver. This is the relevant distinction for comparing in-kind formal care benefits to more flexible benefits that can be spent on informal care.

[11] Consider using changes in minimum wage laws as instruments for the price of formal care. Many formal home care workers earn roughly the minimum wage, so changes in the minimum wage likely shift the supply of formal care. But at the same time, changes in the minimum wage also likely change the opportunity cost of informal care-giving by changing the wage or employment prospects of some potential informal care-givers. This likely shifts the demand for formal care since formal and informal care are closely-related goods.

[12] In principle, care plans or maximum benefit rules could limit the amount of formal care that those receiving the in-kind benefit could consume free of charge and thereby raise the shadow price of formal care above zero. In practice, a variety of evidence suggests that recipients of the traditional in-kind benefit were able to consume as much care as they wished free of charge. See online Appendix Section B.4 for additional details and evidence.

Table 1—Average Hours of Formal Care by Treatment Group

	Near-cash	In-kind	Difference p-value
Overall	6.85	14.19	<0.01
Arkansas	6.29	10.76	<0.01
Florida	7.69	18.60	<0.01
New Jersey	7.01	16.10	<0.01

Notes: Means of formal care consumption in hours per week. *Near-cash* and *In-kind* groups are defined by randomized treatment assignment. p-values test for equality of means. Rows denote different samples.

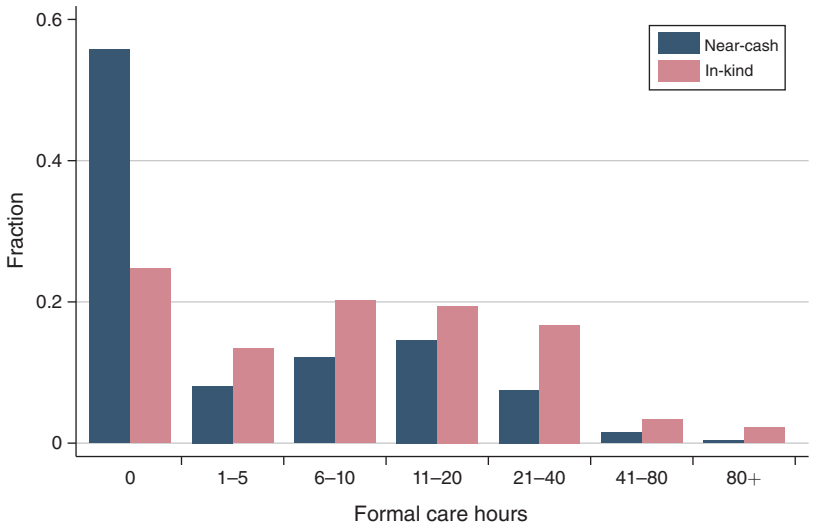

Figure 2. PDFs of Formal Care Consumption by Randomized Benefit Assignment

Notes: Formal care consumption, in hours per week, among participants randomly assigned to the *in-kind* versus *near-cash* benefit. Data from Cash and Counseling follow-up survey.

where p_i is the price of formal care, $Cash_i$ is an indicator of whether the participant was randomized to the near-cash treatment, and X_i includes indicators for sex, education level, race, self-rated health at baseline, living alone at baseline, five-year age bins, and state. The key parameter of interest is β, the effect on formal care consumption of an increase in its net-of-subsidy price. As a starting point, we assume that (ε_i, ν_i) are jointly normal and estimate this system using an instrumental variables Tobit specification.

The first-stage relationship is economically and statistically large. Being randomized to the in-kind benefit decreases the average price of formal care by approximately \$7.70, with a first-stage F-statistic of over 1,100 (see online Appendix Table E.3). The instrumental variables estimate of β is presented in Table 2. It implies that a \$1 increase in the hourly price of formal care reduces consumption by 1.8 hours per week. This corresponds to an elasticity of -1.7 at the sample means. The conclusion that the demand for formal care is highly sensitive to its price holds in each of the three states and is robust to a wide range of alternative assumptions about the distribution of the error terms and benefit limits (see online Appendix Tables E.4

Table 2—The Price Sensitivity of Demand for Formal Care

	(1)	(2)
Price	−1.78	−1.76
	(0.15)	(0.15)
Controls	No	Yes
Mean hours, in-kind	14.19	14.19
Observations	2,440	2,440

Notes: Dependent variable is formal care consumption in hours per week. Specifications are instrumental variables Tobits where formal care hours are censored at zero. *Controls* included in column 2 are indicators for sex, education level, race, self-rated health, five-year age bins, and state. Data are from the Cash and Counseling experiments. Robust standard errors reported.

and E.5). See online Appendix Section C for details and a discussion of the generalizability of the results to other populations and policies of interest.

The estimates imply that in-kind provision has a large moral hazard cost. An individual consuming the average amount of formal care in the in-kind group would consume no formal care without the subsidy and values the care she does receive at just 28 percent of its cost.[13]

IV. Targeting Benefit of In-Kind Provision

In this section, we provide evidence on the two ingredients that determine the targeting benefit of in-kind provision: the distribution of formal care consumption and the link between formal care consumption and the marginal utility of income. We discuss the implications for the targeting benefit of in-kind home care and conclude with evidence on the targeting benefit of Medicaid home care.

A. *Distribution of Formal Care Consumption*

We use data from the National Long Term Care Survey (NLTCS) to estimate the distribution of formal care consumption among the home-care-eligible population. The NLTCS is a nationally representative survey of Americans 65 and older who are eligible for Medicare (see online Appendix Table E.2 for summary statistics). We use the standard eligibility criterion for home care benefits: having at least two ADL limitations. A subset of this population with low enough income and assets is also eligible for Medicaid home care.

Figure 3 shows the distribution of formal care consumption in the home-care-eligible population. Even within this group of people with severe chronic health problems, there is significant heterogeneity in formal care consumption.[14] Sixty-three percent do not consume any formal care. Among those who do there is a long right tail. For that group, the ninety-fifth percentile is around-the-clock care, almost

[13] With $\beta = -1.8$ and no income effects, someone consuming 14 hours of care per week has an equivalent variation of formal care benefits of $54 per week. Medicaid's cost of that care is $192 per week.

[14] The cross-sectional distribution is not a pure measure of risk; it reflects predictable heterogeneity as well as heterogeneity in ex post realizations of risk. In the welfare analysis (Section V), we test robustness to large changes in risk.

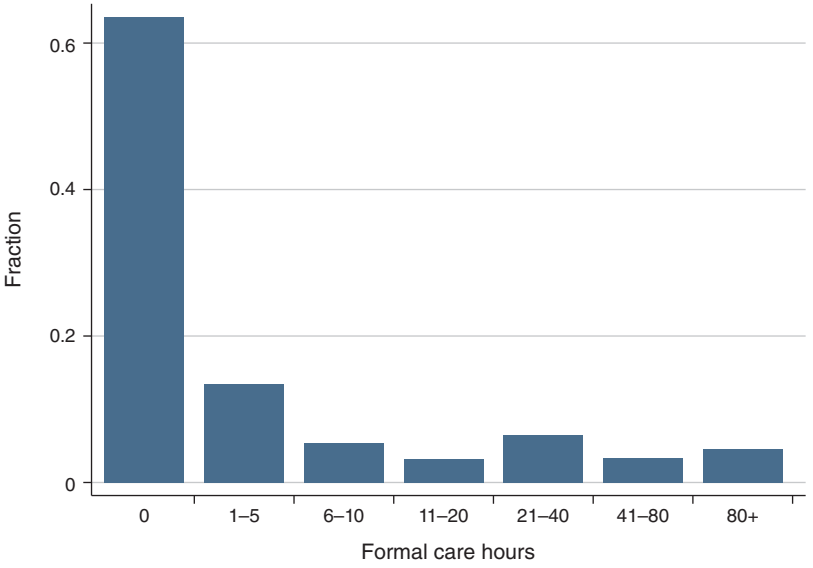

Figure 3. Distribution of Formal Care Consumption in the Benefit-Eligibile Population

Notes: Hours per week of formal care consumption among the non-institutionalized population aged 65 and older with two or more ADL limitations. Data from the 1999 National Long-Term Care Survey. Sixty-three percent do not consume any formal care. Conditional on consuming formal care, median consumption is 10 hours per week, the seventy-fifth percentile is 40 hours per week, the ninetieth percentile is 120 hours per week, and the ninety-fifth and ninety-ninth percentiles are 168 hours per week (around-the-clock care). The standard deviation, σ_{x_k}, is 35 hours per week.

17 times the median among those consuming care. At the average hourly price, that volume of care would cost $131,000 per year.

The significant heterogeneity in formal care consumption implies that in-kind provision has a large targeting effect, sharply concentrating transfers on the small subset of the eligible population with high formal care consumption. The standard deviation of formal care consumption, σ_{x_k}, is 35 hours per week. At the average market price, that implies a standard deviation of annual spending, and so of the ex post marginal benefit of increasing the subsidy rate on formal care, of more than $27,000.

B. *Link between Formal Care Consumption and Marginal Utility*

Both theoretical considerations and empirical evidence suggest a strong link between formal care consumption and marginal utility. In theory, formal care consumption will tend to be positively linked to marginal utility through the budget constraint: greater spending on formal care leaves fewer resources available for non-care consumption.[15] Empirically, private long-term care insurance contracts typically subsidize formal care consumption, and people provide significant informal care and financial support to family members with high formal care consumption.

[15] The idea is that formal care consumption is a poor substitute for "regular," non-care consumption. This is the idea underlying the link between health spending and marginal utility in standard models of health spending risk. See Cutler and Zeckhauser (2000) for a review.

A strong link between formal care consumption and marginal utility implies that in-kind provision of home care would target relatively high-marginal utility states (high $\text{corr}(\hat{\lambda}, x_k)$). Together with the considerable heterogeneity in formal care consumption, a strong link also implies substantial risk within the benefit-eligible population (high $\sigma_{\hat{\lambda}}$). Altogether, this suggests that the marginal targeting benefit of in-kind provision of home care would be large (large $\sigma_{x_k}\sigma_{\hat{\lambda}}\text{corr}(\hat{\lambda}, x_k)$).

C. *Targeting of Medicaid Home Care*

Although the combination of highly-concentrated formal care consumption and a strong link between formal care consumption and marginal utility would imply a large targeting benefit of in-kind provision, the targeting of Medicaid home care depends not only on in-kind provision but also on factors such as awareness of the program and hassles.

Table 3 investigates the targeting of Medicaid home care within the eligible population using nationally representative data from the 1999 NLTCS. The first three rows of the table present estimates of the take-up rate among those eligible for benefits. Differences in the estimates are due to differences in the estimated size of the eligible population. The estimates range from 5 to 19 percent, with 19 percent likely overstating the true rate (see online Appendix Section B.3 for details). Compared to an equal-cost program with complete take-up, low take-up of Medicaid home care increases benefits per recipient by a factor of 5–20. Combining the concentration of benefits from incomplete take-up with that from differences in formal care consumption among those who take up implies a large targeting effect of Medicaid home care. The standard deviation of Medicaid-financed formal care is 27 hours per week.[16]

The next several rows of Table 3 compare the characteristics of those who do versus do not take up benefits among the eligible population, using the *Income eligible, <2 cars* eligibility criteria. People who take up have much greater demand for formal care. If everyone faced a common price, those who take up would be predicted to consume 12 hours per week more formal care on average.[17] Consistent with this, those who take up are sicker (66 versus 46 percent have four or more ADL limitations) and have fewer "prime" potential informal caregivers (67 versus 59 percent are unmarried and 39 versus 29 percent live alone). The correlation of benefits with formal care consumption is 0.62, with number of ADL limitations is 0.13, and with living alone is 0.19.

We turn to investigating targeting in the Cash and Counseling experiments. Unlike take-up of Medicaid home care, the experimental design isolates the effect of in-kind provision. We focus on participants in Arkansas, the only state in which we can calculate each individual's near-cash benefit.

[16] This assumes that all of the formal care consumed by those who take up Medicaid home care is paid for entirely by Medicaid and that all Medicaid home care benefits are in-kind formal care, not cash (cash benefits were rare at the time). That the data lack information on the transfer from Medicaid increases the uncertainty in the calculation but does not obviously bias it toward greater or lesser concentration.

[17] We use our estimated price sensitivity from Section III to predict what each individual's consumption would have been had she faced a price of $18.50 per hour, the maximum price in the data.

Table 3—Targeting of Medicaid Home Care

	Take-up = 0 (1)	Take-up = 1 (2)	Difference *p*-value (3)
Fraction of eligibles who do versus do not take up, under different definitions of eligibility			
Income eligible, <2 cars	0.95	0.05	
Income eligible, no cars	0.90	0.10	
Restrictive income, no cars	0.81	0.19	
Summary statistics			
Level of formal care demand	8.30	20.82	<0.01
Age	80.01	80.82	0.45
Four or more ADLs	0.46	0.66	<0.01
Health fair or poor	0.69	0.78	0.12
Female	0.70	0.72	0.66
Lives alone	0.29	0.39	0.12
Unmarried	0.59	0.67	0.19
Has children	0.75	0.78	0.73
Household income, monthly	847.95	675.56	0.01

Notes: Means for people who did (column 2) versus did not (column 1) take up Medicaid home care. *Difference p-value* tests the equality of means across groups. Take-up rates based on non-institutionalized individuals aged 65 and older with two or more ADL limitations who meet different sets of financial-related eligibility criteria. *Income eligible* is based on the income thresholds each state uses to determine eligibility. *Restrictive income* applies the most stringent (lowest) income limit to all states to try to estimate an upper bound on take-up. Number of cars is an important determinant of eligibility for Medicaid home care. Summary statistics by take-up decision are for those who meet the *Income eligible*, *<2 cars* criteria. This sample has 448 individuals. The level of formal care demand, in hours per week, uses our estimate of price sensitivity to simulate each individual's hours of formal care if she faced a price of $18.50 per hour, the maximum in the data. The alternative to health fair or poor is health good or excellent. Data from the 1999 NLTCS.

Figure 4 shows the distributions of transfers separately for those randomized to the in-kind and near-cash benefits.[18] The in-kind benefit concentrates transfers significantly relative to the near-cash benefit. Transfers to those assigned to the near-cash benefit cluster tightly around the median of $147 per week. Transfers to those assigned to the in-kind benefit are much more dispersed, with a standard deviation more than twice as large and a much greater likelihood of being very large or very small.

Figure 5 shows the extent to which each benefit type concentrates transfers on people with the greatest demand for formal care. For each benefit type, we rank those randomly assigned to that benefit by their formal care consumption. Then we calculate the average transfer, in dollars, received by people at different ranks of the distribution. The in-kind transfers are highly concentrated on those with the greatest demand for formal care. Whereas the average in-kind transfer is $133 per week, individuals between the ninety-first and ninety-fifth percentiles of the formal care distribution receive an average of $350 per week and individuals above the ninety-fifth percentile receive an average of $843 per week, almost seven times the average benefit. The near-cash transfers, by contrast, are roughly constant throughout the formal care distribution, despite being based on individual medical exams.

[18] The near-cash transfers are calculated as the product of care plan hours and the hourly price of care. The in-kind transfers are calculated as the product of hours of formal care used and the hourly price. For both groups, if the individual leaves Medicaid home care we set their transfer to zero. We censor transfers at $600 for the figure but not elsewhere.

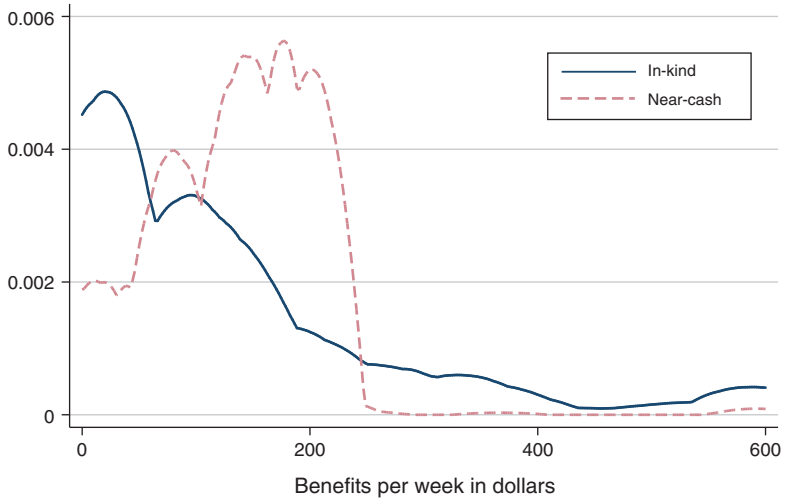

FIGURE 4. TARGETING EFFECTS OF IN-KIND PROVISION ON THE INTENSIVE MARGIN

Notes: Distributions of transfers in the Arkansas Cash and Counseling experiment. Arkansas is the only state for which we observe care plan hours, which we need in order to estimate the near-cash transfer. Transfers are measured in dollar-costs per week at market prices. We scale up the near-cash group's transfers to have the same mean as the in-kind group's in order to isolate differences in the concentration of transfers, not their average size. The average transfer is $133. Groups are based on each individual's randomized assignment. Transfers have been censored at $600 for the figure.

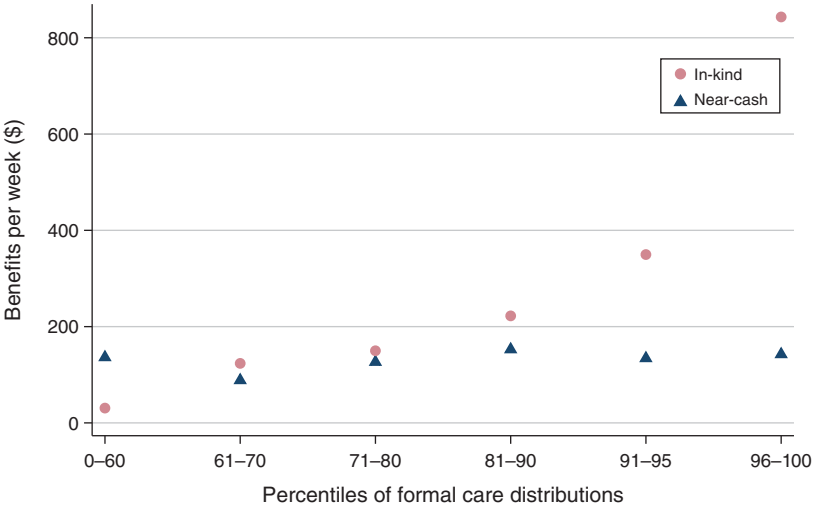

FIGURE 5. TARGETING OF IN-KIND VERSUS NEAR-CASH BENEFITS

Notes: Average transfers, in dollars per week, in the Arkansas Cash and Counseling experiment, separately for those randomized to the in-kind and near-cash benefit. For each benefit type, we rank those randomly assigned to that benefit by their formal care consumption. Then we calculate the average transfer, in dollars, received by people at different ranks of the distribution. Fifty-seven percent of those randomized to near-cash do not consume any formal care.

Online Appendix Section D provides suggestive evidence that in-kind provision concentrates benefits on recipients who are sicker and have fewer informal caregivers than the average recipient as well.

Taken as a whole, these results indicate that in-kind provision sharply concentrates transfers on a small fraction of the eligible population who are sicker, have fewer informal caregivers, and have a greater demand for formal care than the average eligible. To the extent that such recipients tend to have relatively high marginal utility, in-kind provision could have a large targeting benefit.

V. Welfare Effect of In-Kind Provision: Targeting Benefit versus Moral Hazard Cost

This section uses a stylized expected utility model to quantify the net welfare effect of the targeting benefit and moral hazard cost of in-kind home care benefits. As discussed in Section I, the key ingredients for the analysis are the price sensitivity of demand for formal care, the distribution of formal care consumption, and the link between formal care consumption and marginal utility. The first two are readily estimable; the third is provided by the model.

A. *Model*, *Policy Counterfactual*, *and Welfare Measure*

An individual faces risk about her health and her costs of coping with bad health. Together, these determine the level of her demand for formal care. The amount of formal care at which she reaches satiation (i.e., how much she would consume if facing a price of zero) is $\theta \in \mathbb{R}_+$. The θ term is known to be drawn from the distribution $G(\theta)$, but the particular realization of θ is not contractible ex post. Once θ is realized, the individual chooses formal care consumption, F, and non-care consumption, A ("all other goods," the numéraire) to maximize utility subject to a budget constraint that depends on the policy in operation. Indirect utility is

$$v(p, m; \theta) = \max_{A \geq 0, F \geq 0} u\left(A - \frac{(\theta - F)^2}{2\beta}\right) \quad \text{subject to} \quad A + pF = m,$$

where p is the net-of-subsidy price of formal care and m is total after-transfer income, including any cash benefit from the home care program and any transfer from a means-tested program that provides a consumption floor. The corresponding Marshallian demand for formal care is

$$F(p, m; \theta) = \max\left\{0, \min\left\{\frac{m}{p}, \theta - \beta p\right\}\right\}.$$

Note that $\beta \geq 0$ determines the utility cost of consuming levels of care other than the satiation level θ and thereby determines the sensitivity of the demand for formal care to its price.

This utility function is motivated by key evidence from our setting. It produces a simple demand function for formal care that is consistent with some people in

bad health not consuming any formal care, with formal care consumption being sensitive to its price, and with people becoming satiated at finite levels of formal care consumption.[19] It has an intuitive interpretation: utility is decreasing in any unmet, residual care demand, $(\theta - F)$, the size of which is increasing in the level of demand for formal care and decreasing in formal care consumption. This captures the idea that certain health problems are costly for people to cope with on their own. It nests as a special case the widely-used model in which health spending is equivalent to a wealth shock and shares with that model the implication that formal care consumption is linked to marginal utility mainly through the budget constraint. Greater spending on formal care means lower non-care consumption and so greater marginal utility.[20]

We analyze cost-neutral shifts in a mixed in-kind/cash-benefit policy that combines a linear subsidy rate s and a cash benefit b. For any policy, indexed by s, the cash benefit $b(s)$ adjusts to hold fixed total spending on recipients, which is the sum of spending on the subsidy, the cash benefit, and the consumption floor program. Take-up of all benefits is automatic and there are no participation costs. This policy counterfactual isolates the effect of in-kind provision from other sources of incomplete take-up, and it isolates the insurance-moral hazard trade-off of in-kind provision from redistribution between recipients and other parties.

We measure the welfare effect of policy s as its ex ante equivalent variation gain over an equal-cost pure-cash policy, $EV(s)$. Expected (indirect) utility is $EU(s,b) = E(\max\{u(\overline{c}), v(p(s), m + b; \theta)\})$, where $u(\overline{c})$ is utility when relying on the consumption floor. The equivalent variation gain of policy s is the extra income the individual would need over an equal-cost pure-cash policy to be as well off in expected utility as she is under s,

$$EU(0, b(0) + EV(s)) = EU(s, b(s)).$$

B. *Empirical Inputs and Other Parameter Values*

The key empirical inputs are the first two ingredients described in Section I: the price sensitivity of demand for formal care and the distribution of formal care consumption. Our baseline value of the price sensitivity of demand is our main estimate from the Cash and Counseling experiment, $\beta = \left|\hat{\beta}_{C\&C}\right| = 1.8$.[21] This estimate

[19] The most direct evidence of satiation is that among the Cash and Counseling participants for whom we observe care plans, 43 percent consume less care than their care plans entitle them to. Intuitively, satiation might arise from a demand for privacy or space, since home care involves close contact with caregivers in one's home.

[20] As β decreases to zero, demand for formal care becomes less elastic, indirect utility approaches $u(m - p\theta)$, and spending on formal care becomes equivalent to a negative wealth shock: the standard case in the literature on long-term care and health spending risks more generally. Compared to this standard case, our baseline model with $\beta > 0$ implies a weaker link between formal care consumption and marginal utility, which, other things equal, reduces the targeting benefit of in-kind provision. See online Appendix Section E.1 for details.

[21] As described in Section I, the moral hazard cost of in-kind provision depends on the total response of demand to the policy change. Absent income effects on demand for formal care, $\beta_{C\&C}$ is the correct parameter for evaluating any policy that affects the relative price of formal care. With nonzero income effects, $\beta_{C\&C}$ is the right parameter for analyzing policies like those in the Cash and Counseling experiments, which roughly held fixed Medicaid's spending on each participant of the experiments, but not policies with different cash benefits. Cash and Counseling's near-cash benefits were on average greater than those under the policy counterfactual we consider here, which holds fixed total spending on the entire eligible population. With positive income effects on demand for formal care, this estimate will tend to understate slightly the true moral hazard effect of in-kind provision in these policy counterfactuals.

implies that each \$1 increase in the hourly price of formal care reduces formal care consumption by 1.8 hours per week. The high sensitivity of formal care demand to its price means that the moral hazard cost of in-kind provision will be large, especially at high subsidy rates.

Our baseline value of the distribution of formal care consumption is the observed distribution among non-institutionalized individuals age 65 and older who have two or more ADL limitations in the NLTCS. Restricting to people with two or more ADL limitations follows standard practice for Medicaid home care and private long-term care insurance contracts. We use β to convert the observed joint distribution of formal care consumption and formal care prices into the distribution of the level of demand for formal care, $G(\theta)$. Online Appendix Section E.2 contains details of this procedure.

Figure 6 presents our main estimate of the density of the level of demand for formal care, $g(\theta)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits substantial dispersion and has a long right tail. Most of the mass reflects low demand for care; about 56 percent of the θ satiation values are less than 10 hours per week. For those θs, an individual facing the average market price would consume no formal care. But some states have high demand. The ninetieth percentile, for example, is about 37 hours per week. The substantial heterogeneity in demand implies that in-kind provision will concentrate transfers significantly. Together with the model, it also implies substantial heterogeneity in non-care consumption and so in marginal utility. This suggests that the targeting benefit from in-kind provision could be large.

The remaining parameters take standard values. We follow most of the literature on health spending risks and use a constant relative risk aversion utility function, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ (e.g., Brown and Finkelstein 2008, Ameriks et al. 2011). In our model, the argument c is "net consumption," non-care consumption net of any residual coping costs, $c = A - \frac{(\theta - F)^2}{2\beta}$. We follow Brown and Finkelstein (2008) and others in taking as a baseline value a coefficient of relative risk aversion, γ, of 3. Income before transfers is \$15,000 per year. The distribution of before-subsidy prices of formal care is the empirical distribution observed in the NLTCS. If the individual cannot achieve net consumption of at least $\overline{c} = \$5{,}000$ per year, she receives transfers that enable her to reach exactly that living standard. This consumption floor is meant to approximate the combined effects of means-tested government programs like Medicaid and Supplemental Security Income as well as any non-governmental charity care. The higher the consumption floor, the smaller the gains from insurance.[22] The policy counterfactuals hold fixed total spending on recipients at its expected level under a pure in-kind benefit ($s = 1, b = 0$), which is \$7,150.

[22] In many contexts, a sizable fraction of insurance transfers displace means-tested transfers rather than increasing consumption. As a result, greater insurance (a higher subsidy rate in our context) is implicitly taxed by the means-tested program. In the context of long-term care, this implicit taxation of private insurance by means-tested programs is quite large (Brown and Finkelstein 2008). In our baseline specification, the individual relies on the floor in states of the world in which her spending on formal care would otherwise be at least \$10,000, the amount by which income exceeds the consumption floor.

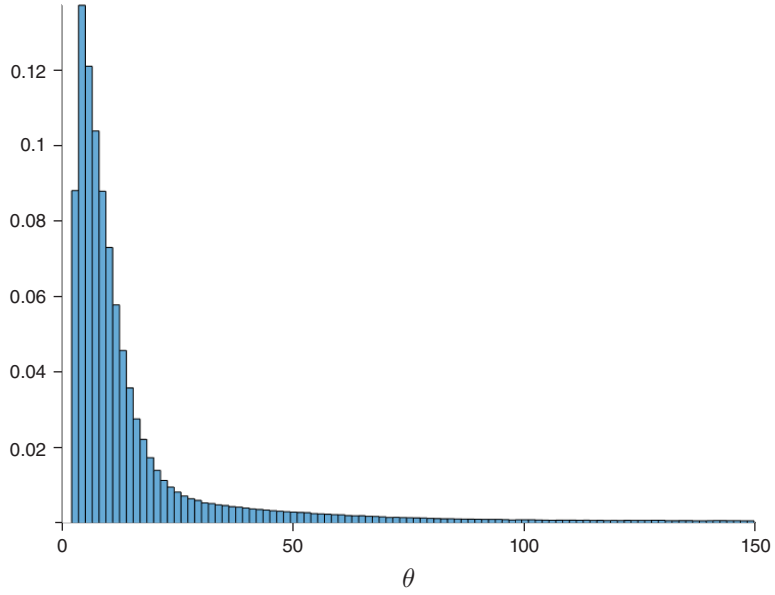

FIGURE 6. DISTRIBUTION OF DEMAND FOR FORMAL CARE

Notes: Simulated distribution of formal care satiation points, θ, in hours per week, among the non-institutionalized population aged 65 and older with two or more ADL limitations. The mean is 16.3 hours per week.

C. *Welfare Effects of In-Kind Provision*

As a benchmark and to get a sense of the extent of the risk within benefit-eligible states, we first calculate the welfare gain from a hypothetical (infeasible) first-best contract that provides state-dependent cash transfers. The equivalent variation gain from this contract is $10,687 (see online Appendix Section E.1 for details).

Figure 7 shows the marginal benefit and marginal cost of a cost-neutral reduction in the price of formal care as a function of the subsidy rate, the quantitative analogue of Figure 1. The marginal benefit far exceeds the marginal cost for subsidy rates up to about 75 percent before falling somewhat below the marginal cost for subsidy rates above 90 percent.[23]

The first column of Table 4 reports statistics about the optimal policy. The optimal subsidy rate is 87 percent, close to that of a pure in-kind program. The optimal subsidy increases welfare substantially. Its equivalent variation gain over the equal-cost cash contract is $6,416. For a cash-benefit contract to achieve the same expected utility, it would have to cost 90 percent more than the optimal contract. Though not optimal, a pure in-kind benefit program with a 100 percent subsidy

[23] The marginal benefit is steeper than the mechanical effect at small subsidy rates because increasing the subsidy increases the set of states in which the individual consumes any formal care. This tends to increase the covariance between marginal utility and formal care consumption. At higher subsidy rates, this effect is dominated by the effect of increasing the subsidy on reducing the variance in marginal utility.

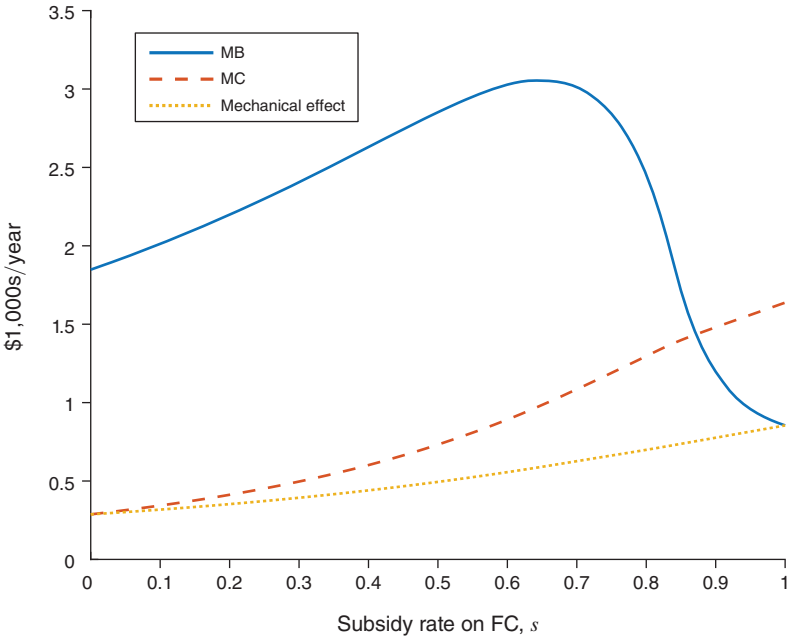

FIGURE 7. MARGINAL BENEFIT AND MARGINAL COST OF DECREASE IN PRICE OF FORMAL CARE

Notes: Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total spending on recipients. $s = 1$ corresponds to a pure in-kind benefit program, a 100 percent subsidy on formal care with no cash benefit. $s = 0$ corresponds to a pure cash benefit program, a 0 percent subsidy on formal care.

TABLE 4—WELFARE ANALYSIS AND KEY ROBUSTNESS TESTS

	Baseline (1)	$\beta = 5$ (2)	Drop $\theta > 50$ (3)	$\theta/2$ (4)	log utility (5)	State-dependent utility (6)
Optimal subsidy, s^*	0.87	0.83	0.57	0.74	0.35	0.60
EV gain over pure-cash policy	\$6,416	\$5,086	\$1,683	\$2,554	\$133	\$1,505
E(ex post value)/E(cost)	0.49	0.59	0.86	0.74	0.80	0.68
corr(marg. util., formal care)	0.89	0.86	0.72	0.88	0.92	0.54

Notes: Subsidy rates are constrained to be no smaller than −0.5 (a 50 percent tax) and no greater than 1.5 (a 150 percent subsidy, under which individuals are paid 50 percent of the market price to consume formal care). *EV gain over pure-cash policy* is the ex ante equivalent variation gain of the optimal policy over an equal-cost pure-cash policy. E(*ex post value*)/E(*cost*) is the ratio of the mean ex post value of the optimal benefit to its mean cost. This is an inverse measure of the distortion cost of the optimal policy. *corr*(*marg. util.*, *formal care*) is the correlation between marginal utility and formal care consumption in the absence of insurance (under a pure-cash policy). This is a measure of how well in-kind provision targets relatively high-marginal utility states. Column 1 corresponds to the baseline assumptions. Column 2 increases the price sensitivity of demand from the baseline estimate of 1.8 to 5. Column 3 truncates the right tail of the risk by dropping all values of θ greater than 50. Column 4 divides every θ by 2, reducing the variance of the risk to one-fourth its baseline value. Column 5 sets the coefficient of relative risk aversion to 1 (log utility), whereas the baseline coefficient of relative risk aversion is 3. Column 6 reduces relative marginal utility in higher-demand states to an extent designed to match the upper end of the most relevant estimates of state dependence in utility, those of Finkelstein, Luttmer, and Notowidigdo (2013).

and no cash benefit also improves substantially on the pure-cash program, with an equivalent variation gain of $5,265.

D. *Robustness and Intuition*

To assess the robustness of the main conclusions and the relative importance of different factors in driving them, we summarize the effects of changes in each of the three key ingredients of the analysis.

Price Sensitivity of Demand for Formal Care.—Our baseline estimate of the price sensitivity of demand implies a large moral hazard cost of in-kind provision. As a result, the optimal contract achieves only 60 percent of the gain from the first-best policy. This shortfall comes from costs along two dimensions. First, recipients overconsume formal care and as a result value the benefit less than its cost in each state ex post. The optimal benefit nearly triples formal care consumption and its mean ex post value is just 49 percent of its cost (column 1 of Table 4). Second, the optimal contract leaves some risk uninsured, since the benefit of insuring it would be more than offset by the moral hazard cost.

Although it is not possible to sign the difference between our estimate of the price sensitivity and the long-run price sensitivity to permanent policies (see online Appendix Section C.3), our finding of a large gain from in-kind provision in the baseline specification makes tests of the robustness of the results to even greater price sensitivity of the most interest. Column 2 of Table 4 shows results based on $\beta = 5$, about three times greater than our main estimate. Although the optimal subsidy and the gain from in-kind provision are smaller than in the baseline specification, they remain large. The optimal subsidy remains large even when $\beta = 10$, over five times greater than our main estimate and greater than seems plausible even in the long run (see online Appendix Table E.8).

Distribution of Formal Care Consumption.—The observed cross-sectional distribution of formal care consumption exhibits substantial heterogeneity with a long right tail. This is the key empirical fact driving the large targeting benefit of in-kind provision. Using the observed distribution to proxy for the (unobservable) counterfactual distribution facing an individual follows much of the literatures on optimal taxation and health care but is limited by the fact that it treats all of the observed heterogeneity as reflecting the result of an exogenous, uncertain process. In reality, some of the observed heterogeneity reflects measurement error, some is predictable, and some is endogenous: people make many choices that affect their future demand for formal care, including investing in their health and relationships with potential caregivers. For these reasons and others, the observed distribution of formal care consumption is an imperfect measure of the risk facing any particular individual.

Columns 3 and 4 of Table 4 show results based on specifications with less variation in the demand for formal care than is implied by the observed distribution of formal care consumption. Column 3 cuts off the right tail of the θ distribution, dropping states with $\theta > 50$. Column 4 scales down the θ distribution, replacing each θ with $\theta/2$ (and thereby reducing the variance of the distribution to one-fourth

of its baseline value). Both changes reduce the optimal subsidy and the gain from in-kind provision, but in both cases the optimal subsidy and the gain from in-kind provision remain large.

Link between Formal Care Consumption and Marginal Utility.—Although our baseline model involves a weaker link between formal care consumption and marginal utility than the benchmark model in which health spending is treated as a wealth shock, in-kind provision is quite effective at targeting states of the world with relatively high marginal utility. Absent insurance (i.e., under a pure-cash policy), the correlation between an individual's marginal utility and formal care consumption is 0.89 (column 1 of Table 4). The strength of this link depends on two sets of factors. One is factors that affect the extent to which greater spending on formal care reduces noncare consumption, including saving and dissaving and informal insurance arrangements. Such factors are most naturally modeled as reducing the dispersion in the distribution of demand for formal care, the effects of which we just discussed.

The other set of factors is features of the utility function, in particular its curvature and any state dependence. Column 5 reduces the coefficient of relative risk aversion to one (log utility). This significantly reduces the extent to which heterogeneity in net consumption translates into heterogeneity in marginal utility and so significantly reduces the value of insurance. Both the optimal subsidy and the gain from in-kind provision are much smaller: 35 percent and $133, respectively. This largely reflects the low cost of home care risk in this specification. The gain from the first-best contract is just $1,215 in this case, 11 percent of its value in the baseline specification. Column 6 considers state-dependence in utility. It reduces relative marginal utility in higher-demand states to an extent designed to match the upper end of the most relevant estimates of state dependence, those of Finkelstein, Luttmer, and Notowidigdo (2013) (see online Appendix Section E.3). This reduces both the extent of risk and how well in-kind provision targets transfers. The worse targeting is revealed by the reduction in the correlation between marginal utility and formal care consumption from 0.89 to 0.54 (columns 1 and 6 of Table 4). As a result, both the optimal subsidy and the welfare gain from in-kind provision are smaller than in the baseline specification. But they remain large in absolute terms: 60 percent and $1,505, respectively.

With one exception, the gain from in-kind home care is large and robust to changes in the key inputs that appear to span the range of plausible values. Online Appendix Section E.4 shows that the results are robust to several other changes as well. The one exception is risk aversion. If risk aversion is low enough, the cost of uninsured risk becomes small enough that even first-best insurance is not that valuable. This points to the key role of the curvature of the utility function in determining the cost of the risk. Provided utility is such that home care risk is important, in-kind provision appears to produce a large welfare gain.

VI. Conclusion

We develop an approach to quantifying a central trade-off of in-kind provision—it can improve targeting at the cost of being less valuable to recipients ex post—and apply it to home care. Despite the ubiquity of in-kind transfers and the centrality of this trade-off to their welfare effects, little is known about the

magnitude of these costs and benefits in many important contexts. We find that in the context of home care, the targeting benefit of in-kind provision appears to exceed its large moral hazard cost. This conclusion is fundamentally driven by the substantial risk within benefit-eligible states of the world.

In focusing on the targeting-distortion trade-off in the stand-alone home care context, we have not explicitly modeled substitution with informal care or nursing homes and we have omitted administrative and take-up costs. Available evidence is suggestive that these considerations would tend to reinforce the net advantage of in-kind provision. In-kind provision likely reduces informal care, which likely increases the labor supply and net tax payments of would-be informal caregivers (Ettner 1995, Van Houtven et al. 2013, Skira 2015), a positive fiscal externality. In-kind provision might also benefit would-be informal caregivers by reducing and insuring their share of the costs of the recipient's bad health. In-kind provision is unlikely to have much effect on usage of nursing homes, given the limited substitution between home care and nursing homes (see Grabowski 2006 for a review). Although cash benefits typically involve lower administrative and take-up costs than in-kind benefits, in home care many of the cash-like benefits that have been implemented in practice involve medical exams, counseling, monitoring, and similar features that likely make them exceptions to this general rule. Whether in-kind provision's targeting benefit could be achieved in less-costly ways is an important question for future research.

Several recent policy reforms and proposals make restrictive in-kind benefits more flexible and cash-like. A major impetus for these reforms is the view that recipients would much prefer equal-cost cash transfers, a view that is consistent with our findings about Medicaid home care. But such reforms also change the distribution of benefits within the eligible population. If a more flexible benefit worsens targeting, the targeting loss should be weighed against the gain from making the benefit more valuable to recipients ex post.

Optimal benefit design is a central policy issue, as many major programs involve in-kind transfers of schooling, housing, food, health care, and other goods. Although home care shares much in common with other important contexts, especially other types of health care, the desirability of in-kind provision is necessarily context-specific. Evaluating the costs and benefits of alternative benefit designs is critically important, and our approach to quantifying this trade-off could prove fruitful in other contexts as well.

REFERENCES

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi. 2016. "Self-Targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124 (2): 371–427.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias. 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *American Economic Review* 102 (4): 1206–40.

Ameriks, John, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh. 2011. "The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Public Care Aversion from Bequest Motives." *Journal of Finance* 66 (2): 519–61.

Arno, Peter S., Carol Levine, and Margaret M. Memmott. 1999. "The Economic Value of Informal Caregiving." *Health Affairs* 18 (2): 182–88.

Barczyk, Daniel, and Matthias Kredler. 2018. "Evaluating Long-Term-Care Policy Options, Taking the Family Seriously." *Review of Economic Studies* 85 (2): 766–809.

Basurto, Pia M., Pascaline Dupas, and Jonathan Robinson. 2017. "Decentralization and Efficiency of Subsidy Targeting: Evidence from Chiefs in Rural Malawi." NBER Working Paper 23383.

Benitez-Silva, Hugo, Moshe Buchinsky, and John Rust. 2004. "How Large Are the Classification Errors in the Social Security Disability Award Process?" NBER Working Paper 10219.

Blackorby, Charles, and David Donaldson. 1988. "Cash Versus Kind, Self-Selection, and Efficient Transfers." *American Economic Review* 78 (4): 691–700.

Brown, Jeffrey R., and Amy Finkelstein. 2008. "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market." *American Economic Review* 98 (3): 1083–102.

Brown, Jeffrey R., and Amy Finkelstein. 2011. "Insuring Long-Term Care in the United States." *Journal of Economic Perspectives* 25 (4): 119–42.

Brown, Randall, S., Barbara Lepidus Carlson, Stacy Dale, Leslie Foster, Barbara Phillips, and Jennifer Schore. 2007. *Cash and Counseling: Improving the Lives of Medicaid Beneficiaries Who Need Personal Care or Home- and Community-Based Services.* Princeton, NJ: Mathematica Policy Research.

Bruce, Neil, and Michael Waldman. 1991. "Transfers in Kind: Why They Can Be Efficient and Nonpaternalistic." *American Economic Review* 81 (5): 1345–51.

Centers for Medicare and Medicaid Services. 2017. National Health Expenditure Accounts. Baltimore: Centers for Medicare and Medicaid Services.

Cohen, Jessica, and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment." *Quarterly Journal of Economics* 125 (1): 1–45.

Cunha, Jesse M., Giacomo De Giorgi, and Seema Jayachandran. 2011. "The Price Effects of Cash Versus In-Kind Transfers." NBER Working Paper 17456.

Currie, Janet. 2006. "The Take-Up of Social Benefits." In *Public Policy and the Income Distribution,* edited by Alan J. Auerbach, David Card, and John M. Quigley, 80–148. New York: Russell Sage Foundation.

Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and in-Kind: Theory Meets the Data." *Journal of Economic Literature* 46 (2): 333–83.

Cutler, David M., and Jonathan Gruber. 1996. "Does Public Insurance Crowd out Private Insurance?" *Quarterly Journal of Economics* 111 (2): 391–430.

Cutler, David M., and Richard J. Zeckhauser. 2000. "The Anatomy of Health Insurance." In *Handbook of Health Economics*, Vol. 1, edited by Anthony J. Culyer and Joseph P. Newhouse, 563–643. Amsterdam: Elsevier.

Da Roit, Barbara, and Blanche Le Bihan. 2010. "Similar and Yet So Different: Cash-for-Care in Six European Countries' Long-Term Care Policies." *Milbank Quarterly* 88 (3): 286–309.

Deshpande, Manasi, and Yue Li. 2017. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." NBER Working Paper 23472.

Doty, Pamela, Kevin J. Mahoney, and Mark Sciegaj. 2010. "New State Strategies to Meet Long-Term Care Needs." *Health Affairs* 29 (1): 49–56.

Ettner, Susan L. 1995. "The Impact of 'Parent Care' on Female Labor Supply Decisions." *Demography* 32 (1): 63–80.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." NBER Working Paper 21308.

Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo. 2013. "What Good Is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption." *Journal of the European Economic Association* 11 (S1): 221–58.

Finkelstein, Amy, and Matthew J. Notowidigdo. 2018. "Take-Up and Targeting: Experimental Evidence from SNAP." NBER Working Paper 24652.

Foster, Leslie, Randall Brown, Barbara Phillips, Jennifer Schore, and Barbara Lepidus Carlson. 2003. "Improving the Quality of Medicaid Personal Assistance through Consumer Direction." *Health Affairs* 22 (S1).

Genworth Financial. 2005. *Cost of Care Survey*. New York: Genworth Financial.

Grabowski, David C. 2006. "The Cost-Effectiveness of Noninstitutional Long-Term Care Services: Review and Synthesis of the Most Recent Evidence." *Medical Care Research and Review* 63 (1): 3–28.

Hoynes, Hilary, and Diane Whitmore Schanzenbach. 2016. "US Food and Nutrition Programs." In *Economics of Means-Tested Transfer Programs in the United States*, Vol. 1, edited by Robert A. Moffitt, 219–300. Chicago: University of Chicago Press.

Jacoby, Hanan G. 1997. "Self-Selection and the Redistributive Impact of In-Kind Transfers: An Econometric Analysis." *Journal of Human Resources* 32 (2): 233–49.

Kleven, Henrik Jacobsen, and Wojciech Kopczuk. 2011. "Transfer Program Complexity and the Take-Up of Social Benefits." *American Economic Journal: Economic Policy* 3 (1): 54–90.

Lepidus Carlson, Barbara, Leslie Foster, Stacy B. Dale, and Randall Brown. 2007. "Effects of Cash and Counseling on Personal Care and Well-Being." *Health Services Research* 42 (1p2): 467–87.

Lieber, Ethan M. J., and Lee M. Lockwood. 2019. "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20180325.

Low, Hamish, and Luigi Pistaferri. 2015. "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off." *American Economic Review* 105 (10): 2986–3029.

Manning, Willard G., and M. Susan Marquis. 1996. "Health Insurance: The Tradeoff between Risk Pooling and Moral Hazard." *Journal of Health Economics* 15 (5): 609–39.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (114): 175–208.

Munro, Alistair. 1992. "Self-Selection and Optimal In-Kind Transfers." *Economic Journal* 102 (414): 1184–96.

Musgrave, Richard A. 1959. *The Theory of Public Finance: A Study in Public Economy*. New York: McGraw-Hill.

National Conference of State Legislatures. 2007. *The State of Long-Term Care.* Washington, DC: National Conference of State Legislatures.

Nichols, Albert L., and Richard J. Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review* 72 (2): 372–77.

Reeder, William J. 1985. "The Benefits and Costs of the Section 8 Existing Housing Program." *Journal of Public Economics* 26 (3): 349–77.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies* 68 (1): 205–29.

Skira, Meghan M. 2015. "Dynamic Wage and Employment Effects of Elder Parent Care." *International Economic Review* 56 (1): 63–93.

Smeeding, Timothy M. 1977. "The Antipoverty Effectiveness of In-Kind Transfers." *The Journal of Human Resources* 12 (3): 360–78.

Van Houtven, Courtney Harold, Norma B. Coe, and Meghan M. Skira. 2013. "The Effect of Informal Care on Work and Wages." *Journal of Health Economics* 32 (1): 240–52.

Zeckhauser, Richard. 1970. "Medical Insurance: A Case Study of the Tradeoff between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory* 2 (1): 10–26.