



Targeting with In-kind Transfers: Evidence from Medicaid Home Care

Ethan M.J. Lieber and Lee M. Lockwood



Targeting with In-kind Transfers: Evidence from Medicaid Home Care

Ethan M.J. Lieber
University of Notre Dame

Lee M. Lockwood
Northwestern University and NBER

January 2017

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
www.mrrc.isr.umich.edu
(734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium through the University of Michigan Retirement Research Center Award #5-RRC08098401-05-00. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Regents of the University of Michigan

Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; Ron Weiser, Ann Arbor; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*

Targeting with In-kind Transfers: Evidence from Medicaid Home Care

Abstract

Many of the most important government programs make transfers in kind as opposed to in cash. Making transfers in kind has the obvious cost that recipients would at least weakly prefer cost-equivalent cash transfers. But making transfers in kind can have benefits as well, including better targeting transfers to desired recipients. In this paper, we exploit large-scale randomized experiments run by three state Medicaid programs to investigate this central tradeoff for in-kind provision. Despite the large distortion from the in-kind provision of formal home care, the benefit from better targeting transfers to high-marginal utility types appears to be even greater. This highlights an important cost of recent policy reforms toward more flexible, cash-like benefits.

Citation

Lieber, Ethan M.J., and Lee M. Lockwood. 2017. "Targeting with In-kind Transfers: Evidence from Medicaid Home Care." Ann Arbor, MI. University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2017-359.
<http://www.mrrc.isr.umich.edu/publications/papers/pdf/wp359.pdf>

Authors' acknowledgements

The authors are grateful to Norma Coe, Gopi Shah Goda, Seema Jayachandran, Brian Melzer, Matt Notowidigdo, Iuliana Pascu, Mike Powell, Diane Schanzenbach, Courtney Van Houtven, and many seminar participants for helpful comments. We thank Jose Carreno, Vishal Kamat, and ZeyuWang for excellent research assistance.

1 Introduction

In-kind transfers are a ubiquitous feature of government programs, private contracts, and charitable programs. The U.S. government spends more than 12 percent of GDP on in-kind health care, child care, and schooling alone (Currie and Gahvari, 2008). The recent Affordable Care Act increased in-kind transfers substantially through expanded Medicaid eligibility and subsidies for health insurance. Making transfers in kind has the obvious cost that recipients would at least weakly prefer cost-equivalent cash transfers. But making transfers in kind can have benefits as well, including better targeting transfers to desired recipients (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988).

We analyze this central tradeoff of in-kind provision in the context of long term care. Although long-term care is one of the largest financial risks people face over the life cycle, much of it is not insured. We investigate the importance of a particular barrier to insuring this risk: costly state verification. In particular, we investigate the extent to which hard-to-verify heterogeneity in the cost of coping with bad health constrains risk sharing and shapes the optimal benefit design.

To see the potential importance of costly state verification in long-term care, consider the following example. Two 80-year-olds have the same severe chronic health problems that make difficult even simple activities like eating and bathing. Each faces two possibilities: receive a significant amount of assistance at home or go to a nursing home—something both would rather avoid. The key difference between the two people is that one has a smaller cost of coping with bad health: One has a healthy spouse who is willing and able to provide informal care, whereas the other must purchase costly formal home care (eight hours per day would cost approximately \$40,000 per year) if he is to avoid going to a nursing home. This difference in the cost of coping with bad health means that the same health problems have vastly different effects on the marginal utility of wealth of the two individuals.

To the extent that ex-post differences in the cost of coping with bad health are uncertain ex ante, the first-best insurance contract makes larger transfers to people with larger coping costs. But such transfers are infeasible in practice if coping costs cannot be contracted on. It may be difficult, for example, to contract on the preferences of an individual and her family about privacy and close contact. In the second-best case in which coping costs cannot be verified at reasonable cost, insurance programs must offer the same benefit to everyone. Because a cash benefit can be expected to attract people regardless of their coping costs, it cannot implement the (ex-ante) valuable transfers from people with smaller coping costs to people with larger coping costs. This is a major failure of insurance.

This raises the possibility that non-cash benefits might improve risk sharing. A natural candidate is an in-kind benefit of formal home care, which is likely to be more valuable to people with larger coping costs. To the extent that people with smaller coping costs are less likely to take up the home care benefit, in-kind home care is targeted toward people with greater coping costs and, presumably, higher marginal utility. Even with in-kind benefits, however, information frictions limit the extent of risk sharing since an attractive home care benefit may appeal even to people with relatively low coping costs. As a result, there is a fundamental tradeoff between providing benefits that are more valuable to recipients (for which less restrictive cash-like benefits are best) and providing benefits that better target transfers to higher-marginal utility types (for which more restrictive in-kind benefits are best).

In principle, both the costs and benefits of providing transfers in kind could be large. Empirical work is therefore essential for evaluating these policies. Yet little empirical work has investigated the costs and benefits of providing transfers in kind in a unified framework.

In this paper, we develop a general approach for estimating the targeting benefit and distortion cost of in-kind provision that shares much in common with the literature on optimal taxation. We use this approach to investigate the consequences of providing formal home care benefits in kind. The theory highlights the central importance of two key factors in determining the desirability of in-kind transfers: the sensitivity of the demand for the good to the composition of benefits and the distribution of the level of demand for the good.

To estimate the sensitivity of the demand for formal care to the composition of benefits, we take advantage of large-scale experiments run by three state Medicaid programs known as the Cash and Counseling demonstrations. These experiments randomized a subset of Medicaid in-kind home care participants to a (near-)cash benefit. Existing research on the Cash and Counseling experiments has focused on characterizing the effects of in-kind versus cash benefits on various outcomes, such as recipients' satisfaction, health, and healthcare consumption (see Brown et al., 2007, for a review). Key findings include that people randomized to the traditional Medicaid in-kind home care benefit report lower satisfaction and slightly worse health. These findings have led to calls for more flexible, cash-like benefits. Yet to our knowledge, no research has attempted to quantify the net welfare effects of in-kind provision of home care benefits.

The Cash and Counseling experiments are especially useful for evaluating the consequences of in-kind provision for two reasons. First, estimates of the demand for formal care based on non-experimental data face a particularly difficult identification problem: Many factors that shift the supply of formal care also shift the supply of *informal* care—and thereby shift the demand for formal care, causing simultaneity bias. Second, the price variation generated

by the experiments spans the range of interest for policy, from the market price to a zero price. Our results reveal that the demand for formal home care is highly sensitive to the composition of benefits. On average, people randomized to the in-kind transfer consumed twice as much formal care as those randomized to the cash transfer. We estimate that a one dollar increase in the price of formal care reduces consumption by 1.8 hours per week. This price sensitivity implies that providing someone 14 hours of formal home care per week (the average in the Cash and Counseling demonstration states) causes an ex-post deadweight loss of 72 percent of Medicaid’s spending on the individual’s care. In-kind provision of formal care apparently causes a large and costly distortion.

We estimate the second key factor determining the desirability of in-kind transfers of formal care—the distribution of demand for formal care—using the National Long-Term Care Survey. The data reveal significant heterogeneity in the level of demand for formal care. Even among people with severe chronic health problems, most do not consume any formal care, mainly because they rely on informal care instead (Barczyk and Kredler, 2016). Among those who do consume formal care, the distribution of care consumption features a long right tail; some people consume large amounts of care at significant out-of-pocket cost. Moreover, much of the heterogeneity in formal care is uncorrelated with any observables in the data, including the main candidates one might think to use as “tags” in order to direct benefits to high-marginal utility groups (Akerlof, 1978).

Using our estimates of the sensitivity and distribution of demand for formal care, we analyze the welfare consequences of alternative home care benefit designs. First, we take as given the standard eligibility criteria for home care benefits and compare alternative benefit designs, from a pure in-kind benefit (a 100 percent subsidy rate on formal care) to a pure cash policy (no subsidy on formal care). We find that across a wide range of assumptions, the optimal subsidy on formal care is significantly greater than zero and is often fairly close to one (under which the out-of-pocket price for formal care is zero). Second, we estimate the value of using tags—conditioning benefits on individuals’ verifiable characteristics such as objective measures of health. We find that even tags based on characteristics that are highly correlated with formal care consumption have little benefit, since much of the relevant heterogeneity occurs within groups of people with the same observable characteristics. Together, these results show that despite the large distortion caused by the optimal formal care subsidy, welfare is significantly greater under the subsidy than it is under alternative cash-benefit programs, including those that use tags to offer different cash benefits to different groups. The desirability of providing formal care in kind arises from the significant heterogeneity in the demand for formal care, which in the absence of a large formal care subsidy translates into significant heterogeneity in non-care consumption and marginal utility.

These results have important policy implications. The current equilibrium—in which much risk is uninsured, distortions are large, and the means-tested Medicaid program is the primary payer—has many obvious drawbacks. Many recent policy reforms seek to improve matters by making benefits more flexible or cash-like.¹ Yet to our knowledge, no prior research has attempted to measure the costs and benefits of these reforms in a unified framework. Our analysis fills this gap. On the one hand, our finding that formal care consumption is highly sensitive to the composition of benefits implies, consistent with the arguments of advocates for more cash-like benefits, that many recipients would be significantly better off *ex post* with a cost-equivalent cash transfer. It also implies that a large “moral hazard tax” plagues most long-term care insurance contracts—an under-appreciated difficulty facing a market with many other challenges besides—and raises the effective loads to consumers above existing estimates (e.g., Brown and Finkelstein, 2007; Friedberg et al., 2014). On the other hand, our analysis suggests that the large cost of providing home care in kind might be an unfortunate but necessary cost of insuring the risk from chronic health problems. The cost of coping with chronic health problems appears to vary greatly, even within groups of people with the same observable characteristics. As a result, under a wide range of assumptions it is optimal to have a significant in-kind component of home care benefits. This finding raises concerns about recent reforms that make long-term care benefits more flexible and cash-like.

A central contribution of this paper is to develop an approach for estimating the targeting benefit and distortion cost of in-kind provision. We do this by adapting a standard approach in the literature on optimal taxation to the context of in-kind benefits. This provides a link between the theoretical and empirical literatures on in-kind benefits, which, as Currie and Gahvari (2008) note, have been largely disconnected so far.² The theoretical literature on in-kind benefits investigates a variety of potential benefits of in-kind provision in addition to targeting, including increasing the efficiency of the tax system (Munro, 1992), reducing moral hazard in the context of the Samaritan’s Dilemma (Bruce and Waldman, 1991), internalizing externalities, indulging paternalistic preferences, shifting prices in a desirable way (“pecuniary effects,” Cunha et al., 2011), mitigating asymmetric information problems, solving political economy problems, and redistributing resources within households or fami-

¹Several European countries, including Germany, France, and Italy, have major cash-benefit programs (Da Roit and Le Bihan, 2010). In the United States, fifteen states created Cash and Counseling programs by 2009 (Doty et al., 2010), and early versions of the bill that became the Affordable Care Act included a long-term care insurance program that would have paid cash benefits. This program, known as the CLASS (Community Living Assistance Services and Supports) Act, was eventually repealed due to concerns about its budgetary sustainability.

²A canonical in-kind transfer is equivalent to a non-linear price subsidy, under which purchases of the good are subsidized at a 100 percent rate up to a fixed quantity limit beyond which further purchases are not subsidized. The key feature of an in-kind transfer is that it changes the relative price of the transferred good. This feature is shared by many programs that make what appear to be cash transfers as well. For example, health insurance reimbursements (in cash) for health care expenses incurred depend on actions taken by the insured (consuming care).

lies (see Currie and Gahvari, 2008, for a review). Some of these other potential advantages of in-kind provision may be important in the context of home care as well, especially tax system efficiency (since providing informal care appears to reduce market work, e.g., Ettner, 1995) and perhaps the Samaritan's Dilemma (providing care to one's elderly parents may make it more likely that one will rely on means-tested transfers in the future). These other potential benefits are an interesting topic for future research. The empirical literature on in-kind benefits is mostly focused on the consumption distortion cost of providing benefits in kind, especially in the case of food stamps. Moffitt (1989), Whitmore (2002), and Hoynes and Whitmore Schanzenbach (2009), for example, all find that providing food in kind (via food stamps) has relatively small effects on recipients' choices relative to cost-equivalent cash transfers. This is because the benefit levels are lower than many recipients' food purchases and because some possibility of resale exists. The apparent similarity of food stamps to cash reduces the likelihood that they have important targeting effects.³ Our hope is that the approach we develop in this paper proves fruitful in other contexts as well.

2 Theory

This section presents a simple model of in-kind benefits that focuses on the targeting benefit and distortion cost of in-kind provision. It combines features from the welfare analysis of a tax or subsidy, where efficiency costs are central, with features from optimal income taxation and the welfare analysis of social insurance, where redistribution across types is central.

The key feature of in-kind provision is that the size of the transfer an individual receives depends on his or her consumption of the good in question. One can view an in-kind benefit program as providing a cash benefit while at the same time imposing a restriction on recipients that they must consume at least a certain amount of the good in question. As Nichols and Zeckhauser (1982) emphasize, imposing restrictions on recipients can improve the targeting of benefits to desired recipients who cannot otherwise be distinguished from would-be "mimics," if meeting the restriction is more costly for mimics than for desired recipients. Imposing such a restriction relaxes the incentive compatibility constraints on mimics' participation and thereby allows the program to make greater transfers to desired recipients.

An in-kind benefit can be modeled as a (potentially non-linear) price subsidy. Many in-kind benefit programs, such as food stamps, offer individuals up to a fixed quantity of the good at no charge. When resale is not possible, this has the same effect on a participating

³Food stamps may have important targeting effects even if they do not distort consumption if stigma or other factors affect take up.

individual’s budget constraint as a non-linear price subsidy of 100 percent on units up to the benefit limit and 0 percent on units above the limit.⁴ In this section we focus on the case of a subsidy program with no quantity limit. We do this both for simplicity of exposition and because in many states, including the states that ran the Cash and Counseling experiments, the Medicaid home care program does not appear to have binding benefit limits in practice. The results are easily extended to cases with benefit limits.

The key considerations for in-kind provision can be seen in Figure 1. Figure 1 shows the values (in terms of equivalent variations) and efficiency costs of a price subsidy on good X for each of two consumers with different levels of demand for X . The price subsidy is worth less to each consumer than it costs the government or insurance company to provide due to the induced consumption distortion. The size of this distortion is increasing in the compensated own-price elasticity of demand. This is entirely standard. Less-commonly emphasized is the key role of heterogeneity in the level of demand for X in determining both the distributional consequences of subsidizing X and the share of total spending on the subsidy that is costed away by distortions. The same price subsidy is more valuable to people who consume more of the subsidized good, so, relative to a cost-equivalent cash benefit, the subsidy redistributes toward people who consume more X from people who consume less X . In addition, the share of the program’s spending on the subsidy that benefits the recipient as opposed to being costed away by distorting consumption is increasing in the recipient’s level of demand for the good. As a result, the extent to which consumption distortions reduce the desirability of in-kind provision depends not only on the price sensitivity of demand but also on the distribution of the level of demand across types.

2.1 The benefit program and its budget constraint

Consider the problem of designing a mandatory benefit program for a population of ex-ante identical individuals whose ex-post distribution of types is $F(\theta)$.⁵ An individual’s type, $\theta_i \sim F(\theta)$, embeds all of the individual’s characteristics that are relevant for determining the costs and benefits of alternative benefit designs, including any relevant heterogeneity in preferences and budget constraints. The planner knows the distribution of types, $F(\theta)$, but

⁴The nature of resale opportunities, if any, is an important determinant of the effects of in-kind benefit programs. In the case of home care benefits, resale is impossible. In the case of food stamps, by contrast, resale markets are an important feature of the environment. Whitmore (2002) presents survey evidence that food stamps trade at about 65 percent of their face value in the resale market.

⁵That the program is mandatory is not essential. It just simplifies the exposition by eliminating take-up decisions.

cannot verify any single individual's type.^{6,7}

Consider an idealized in-kind benefit program that potentially combines two elements: a cash benefit, b , and a linear subsidy on good X , σ . The cash benefit and subsidy rate are common across all eligible individuals and are automatic in the sense that there are no take-up decisions; all eligible individuals receive the cash benefit and are subsidized on their purchases of good X . Two special cases of this combined cash-plus-subsidy program are a pure cash-benefit program ($b > 0, \sigma = 0$) and a pure subsidy program with no cash benefit ($b = 0, \sigma > 0$). A pure in-kind benefit program like Medicaid home care has a zero cash component and a full subsidy, ($b = 0, \sigma = 1$).

Average per-eligible spending on the program, B , is divided between funding the cash benefit, b , and the subsidy on X , σ :

$$\int_{\Theta} (b + (\sigma p_X(\sigma))x_X(\sigma; \theta)) f(\theta) d\theta = B,$$

where $p_X(\sigma)$ is the subsidy-exclusive price of X (the sellers' price) and $x_X(\sigma; \theta)$ is the consumption of X by type θ as a function of the subsidy rate.

2.2 Analysis of a budget-neutral shift toward in-kind benefits

This section analyzes a marginal shift in benefits toward in-kind benefits. This shift involves marginally increasing the subsidy rate, σ , and at the same time decreasing the cash benefit in order to maintain the same program budget.

For simplicity, suppose that the supply of every good is perfectly elastic. In this case, an increase in the subsidy reduces buyers' after-subsidy price of X one-for-one (no incidence on supply), $p_X(\sigma) = (1 - \sigma)p_X^0$, and has no effect on the prices of goods other than X , $p_i(\sigma) = p_i^0$ for $i \neq X$, where p_i^0 is the price of good i without any benefit program.

⁶We focus on the problem of designing benefits for a population of ex-post heterogeneous individuals who cannot be verifiably distinguished. This population could be a sub-population of a broader population, where the sub-population is distinguished from the broader population by the values of some verifiable characteristics. Differences in verifiable characteristics allow the program to "partition" the population into sub-groups and treat these groups differently. Groups of individuals who can be verifiably distinguished can be given different benefits directly; the planner need not resort to imposing restrictions on recipients in order to redistribute across these groups. In the extreme case in which all of the heterogeneity in marginal utility occurs across rather than within groups, pure cash-benefit contracts can achieve the first-best allocation. In the opposite extreme in which all of the heterogeneity in marginal utility occurs within rather than across groups, pure cash-benefit contracts can provide no insurance within the eligible population.

⁷For simplicity, we ignore any second-best considerations that might arise from the interaction between the program and other distortions in the economy. We discuss such considerations in Section 6 and the conclusion.

Marginally increasing the subsidy rate while at the same time decreasing the cash benefit in order to maintain the same program budget implies the following change in the cash benefit:

$$\begin{aligned}\frac{\partial b(\sigma, B)}{\partial \sigma} &= - \int_{\Theta} \left[x_X(\sigma; \theta) p_X^0 + (\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right] f(\theta) d\theta \\ &= - \left[E_{\Theta} (x_X(\sigma; \theta) p_X^0) + E_{\Theta} \left((\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right) \right].\end{aligned}$$

The cash benefit must fall by the increase in average per-eligible spending on the in-kind benefit (subsidy). Average spending on the subsidy is the sum of two terms: (i) the mechanical increase in spending on the subsidy due to the increase in the subsidy rate, holding fixed each type's consumption of X , $E_{\Theta} (x_X(\sigma; \theta) p_X^0)$ (“mechanical effect”); and (ii) the increase in spending on the subsidy due to the induced change in consumption of X in response to the shift in program benefits, $E_{\Theta} \left((\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right)$ (“behavioral effect”).⁸

2.2.1 The net ex-post value for each type of a shift toward in-kind provision

For each type, the net ex-post value of the change in the program is the benefit of the increase in the subsidy on X (i.e., the benefit from the reduction in the after-subsidy price of X) less the cost of the reduction in the cash benefit. A marginal increase in the subsidy rate on X reduces the after-subsidy price of X by p_X^0 . The value (in units of income) of this reduction in the price of X to an individual of type θ is, by the envelope theorem (Roy's identity),

$$\frac{\frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial p_X(\sigma)} \frac{dp_X(\sigma)}{d\sigma}}{\frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial m}} = x_X(\sigma; \theta) p_X^0,$$

where $v(p(\sigma), m(\sigma, B); \theta)$ is the indirect utility function of type θ and $m(\sigma, B) = m^0 + b(\sigma, B)$ is benefit-inclusive income. This benefit from a lower after-subsidy price of X must be weighed against the reduction in the cash benefit required to hold fixed total spending on the program. Combining these two elements gives the net value (in units of income) of a

⁸The “behavioral effect” can be positive or negative, though in most cases it will be positive. It embeds the income effects from the reduction in cash benefits, which tend to reduce the consumption of X (provided X is normal), and substitution and income effects from the reduction in the after-subsidy price of X , which tend to increase consumption of X . A shift toward in-kind provision increases average consumption of X unless income effects of demand for X are much larger among those who lose from the shift than among those who gain.

budget-neutral marginal shift toward in-kind benefits of

$$\begin{aligned}
\frac{dV(\sigma; \theta)}{d\sigma} &\equiv \frac{\frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma}}{\frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial m}} = \frac{\frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial p_X} \frac{dp_X(\sigma)}{d\sigma} + \frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial m} \frac{\partial b(\sigma, B)}{\partial \sigma}}{\frac{\partial v(p(\sigma), m(\sigma, B); \theta)}{\partial m}} \\
&= x_X(\sigma; \theta) p_X^0 - \int_{\Theta} \left[x_X(\sigma; \theta) p_X^0 + (\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right] f(\theta) d\theta \\
&= [x_X(\sigma; \theta) p_X^0 - E_{\Theta}(x_X(\sigma; \theta) p_X^0)] - (\sigma p_X^0) E_{\Theta} \left(\frac{dx_X(\sigma; \theta)}{d\sigma} \right). \quad (1)
\end{aligned}$$

The marginal net value for an individual of type θ of a budget-neutral marginal shift in benefits toward in-kind benefits is the net benefit of the resulting redistribution to his type (redistribution benefit), $[x_X(\sigma; \theta) p_X^0 - E_{\Theta}(x_X(\sigma; \theta) p_X^0)]$, which is greater for types with greater levels of demand for X , less the average marginal distortion cost from the induced change in consumption of X (distortion cost), $(\sigma p_X^0) E_{\Theta} \left(\frac{dx_X(\sigma; \theta)}{d\sigma} \right)$.

Equation 1 shows that the shift toward in-kind provision has two key effects. It redistributes toward people with above-average demand for the good, and it distorts consumption of the good. The extent to which a particular type gains from a marginal shift toward greater in-kind provision is increasing in that type's level of consumption of the good and decreasing in the average sensitivity of the demand for the good in the population. Absent heterogeneity in the consumption of X , all types lose from the shift toward in-kind provision. More generally, a marginal shift toward in-kind provision makes *all* types worse off (a Pareto worsening) if heterogeneity in the consumption of X is small enough relative to the distortion caused by the subsidy that any redistribution gain enjoyed by high- X types is more than offset by the reduction in the cash benefit due to the distortion.

2.2.2 The net ex-ante value of a shift toward in-kind provision

Ex-ante expected utility is

$$\max_{\sigma} EU(\sigma) = \int_{\Theta} v(p(\sigma), m(\sigma, B); \theta) f(\theta) d\theta.$$

The total derivative of expected utility with respect to the in-kind component σ (adjusting the cash component b in order to hold fixed total program spending) is:

$$\begin{aligned}
\frac{dEU(\sigma)}{d\sigma} &= \int_{\Theta} \frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma} f(\theta) d\theta = \int_{\Theta} \lambda(\sigma; \theta) \frac{dV(\sigma; \theta)}{d\sigma} f(\theta) d\theta = E_{\Theta} \left(\lambda(\sigma; \theta) \frac{dV(\sigma; \theta)}{d\sigma} \right) \\
&= Cov_{\Theta} [\lambda(\sigma; \theta), x_X(\sigma; \theta) p_X^0] - (\sigma p_X^0) E_{\Theta}(\lambda(\sigma; \theta)) E_{\Theta} \left(\frac{dx_X(\sigma; \theta)}{d\sigma} \right), \quad (2)
\end{aligned}$$

where $\lambda(\sigma; \theta)$ is the marginal utility of income.

Equation 2 shows the key roles of heterogeneity in the level of demand for X and the sensitivity of the demand for X to the composition of benefits in determining the welfare consequences of in-kind provision. The extent of heterogeneity in the demand for X and the extent to which this heterogeneity is correlated with marginal utility together determine the targeting benefit of in-kind provision. The greater is the covariance across types in marginal utility and the demand for X , the greater is the targeting benefit of in-kind provision. The sensitivity of the demand for X to the composition of benefits determines the distortion cost of in-kind provision. The greater is the sensitivity of the demand for X to the composition of benefits, the greater is the distortion cost of in-kind provision.⁹

3 Long-Term Care Risk and Financing and the Design of Home Care Benefits

Chronic health problems are the source of one of the most important risks people face over the life cycle. Most people will at some point develop functional limitations and require significant, time-intensive assistance with activities such as bathing, eating, dressing, and managing their household. Roughly 15 percent of Americans over age 50 have at least one functional limitation and at least one person helping them as a result.¹⁰ People deal with their functional limitations in a wide variety of ways. Among people receiving help, the vast majority (87 percent) live in the community; only 13 percent live in nursing facilities, mostly nursing homes. Among those living in the community, most (72 percent) receive only informal care from family and friends, 10 percent receive a mix of formal and informal care, and only 4 percent rely exclusively on formal home care. (We will refer to formal home care as “formal care” and to informal home care as “informal care.”) Of all hours care, 64 percent are informal care, 10 percent are formal home care, and 27 percent are from care-giving facilities. The costs of such care are quite large. In 2008, total spending on formal long-term care was \$203 billion, with one-third going to home and community based services

⁹Appendix A analyzes the optimal mix of in-kind and cash benefits. Absent heterogeneity in the demand for X , the optimal policy is a pure cash benefit with no subsidy on X , ($b = B, \sigma = 0$). Absent any consumption distortion, the in-kind benefit simply redistributes resources across different types (as defined by their level of demand for X), at no efficiency cost. In this case, the optimal policy eliminates the covariance between marginal utility and the demand for X . If the demand for X is at least somewhat elastic, by contrast, the optimal policy trades off the insurance benefit of increasing in-kind provision against the distortion cost. In most cases it will stop short of eliminating the covariance between marginal utility and the demand for X , since at the margin there would be only a distortion cost and no targeting benefit.

¹⁰This and the other statistics about care-giving arrangements in this paragraph come from Barczyk and Kredler (2016).

(Centers for Medicare and Medicaid Services, 2010). The cost of informal care is estimated to be no less than 60 percent of, and is often thought to exceed, the cost of formal care (e.g., Arno et al., 1999).

Although the costs of dealing with chronic health problems are highly uncertain and sometimes quite large, much of these costs are uninsured. In many countries, universal health insurance programs (like Medicare in the U.S.) do not cover long-term care. Yet markets for private long-term care insurance are uniformly small. In the U.S., just 10 percent of people 65 and older own private long-term care insurance, and coverage rates are similarly low in other countries. The large costs of long-term care together with the low rates of insurance leads to a situation in which means-tested government programs play a major role in financing long-term care costs. In the U.S., for example, the means-tested Medicaid program covers at least part of the costs of 70 percent of nursing home residents at any point in time (Kaiser Commission on Medicaid and the Uninsured, 2013).

Medicaid also plays a major role in financing home care. In most states, Medicaid provides home care primarily through two programs: the Medicaid Title XIX PCS optional State plan and the Medicaid 1915(c) HCBS waiver program. Eligibility is based primarily upon financial means tests and an assessment of the individual’s “need” for home care based on her health.¹¹ Medicaid home care programs have grown rapidly in recent years. In 1999, 1.9 million people received home care through Medicaid; by 2013, that number had increased to nearly 3 million. In addition to the growing number of participants, the fraction of Medicaid long-term care dollars that go to home care has risen from 18 percent in 1995 to 51 percent in 2014 (Kaiser Commission on Medicaid and the Uninsured, 2016).

Traditionally, government and private long-term care insurance programs have provided in-kind formal care benefits (including subsidies on formal care) to people who meet health- and sometimes financial-related eligibility criteria.¹² But in recognition of the importance of informal care and other means of dealing with chronic health problems, many programs have shifted toward more flexible, cash-like benefits. These programs tend to allow people to spend their benefits on a much wider array of options than just the traditional formal care, often including paying family and friends for informal care, and at the extreme including

¹¹Medicaid is financed jointly by the federal and state governments, so Medicaid policies vary somewhat across states. For the elderly, the means tests for Medicaid home care are often less restrictive than those for general Medicaid coverage. The majority of states provide coverage for individuals with incomes up to 300 percent of the monthly Supplemental Security Income amount (LeBlanc et al., 2001). The amount of Medicaid home care for which an individual qualifies is determined by a medical exam. The applicant’s health care provider must submit a care plan that details the services deemed appropriate based on the applicant’s health status. Summaries of Medicaid-provided home care services are available in LeBlanc et al. (2001) and Kaiser Commission on Medicaid and the Uninsured (2011).

¹²Many programs, for example, offer benefits to people who have limitations with at least two activities of daily living.

anything that a recipient wishes to purchase, whether care-related or not. 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). In the U.S., many state Medicaid home care programs provide flexible, cash-like benefits, including 15 states with “Cash and Counseling” programs (Doty et al., 2010). Early versions of the Affordable Care Act included a voluntary, long-term care insurance program in which individuals meeting health-related eligibility criteria would have received a daily cash benefit.

The shift toward more flexible, cash-like benefits is motivated by the perception, supported by a range of evidence, that many people with chronic health problems would much prefer a cash transfer to a cost-equivalent in-kind transfer of formal care. An important milestone in the debate about more- vs. less- flexible benefits was the Cash and Counseling demonstrations. These were large-scale experiments run by three state Medicaid programs around the year 2000. The experiments randomized participants between the traditional in-kind formal home care benefit and an alternative more flexible, cash-like benefit that allowed recipients to spend the benefit on a wide range of care-related goods and services, including paying relatives or friends for providing care.

The main goal of the Cash and Counseling experiments was to test whether recipients could effectively manage their cash benefits and receive “enough” care. The results were almost uniformly positive. Members of the cash-benefit treatment group reported greater satisfaction with their care (Foster et al., 2003) and had similar health outcomes; if anything, their health outcomes appeared to be slightly better than those of the traditional in-kind benefit group (Lepidus Carlson et al., 2007). In the official final report on the experiments, Brown et al. (2007) conclude that the cash transfer has overwhelmingly positive effects on recipients. The large body of work evaluating the Cash and Counseling experiments demonstrates quite conclusively that more flexible, cash-like benefits can be much more valuable to recipients than traditional in-kind benefits. This points to an important cost of providing home care benefits in kind: The consumption distortion from in-kind provision appears to be large in this context.

Although much is known about the consumption distortion from in-kind provision of Medicaid home care benefits, little is known about the potential benefits of in-kind provision, whether for Medicaid home care or for other programs more generally (Currie and Gahvari, 2008). In-kind provision can have a variety of effects in addition to distorting consumption, including changing the distribution of benefits among those eligible for benefits. This could significantly increase targeting efficiency, especially in the many contexts in which state verification is particularly costly. Yet little is known about the likely magnitude of such benefits.

This is the gap we aim to fill in this paper.

In the context of long-term care, the diversity of ways in which people cope with chronic health problems and the large role of informal care may raise important challenges for insuring this risk. Hard-to-verify heterogeneity in the costs of coping with chronic health problems—which may arise from hard-to-verify differences in health status or in would-be care-givers’ opportunity costs of providing informal care—seem likely to cause heterogeneity in marginal utility and in the demand for formal care. In a second-best setting in which this heterogeneity is not verifiable, insurance programs cannot make different cash transfers to people on the basis of their coping costs. Any redistribution from people with lower to higher coping costs must instead come through changes in prices. A potentially worthwhile distortion might be to subsidize formal care. Everyone in the same market faces the same price of formal care, which makes formal care consumption a potentially good “meter” of the total costs of bad health (of which formal care is only a part). A formal care subsidy therefore tends to redistribute to people with larger coping costs at the cost of distorting consumption. As discussed in Section 2, two of the main determinants of the desirability of such a subsidy are the distribution of formal care consumption and the sensitivity of formal care demand to the composition of benefits. We now turn to estimating these objects.

4 Sensitivity of the Demand for Formal Care to the Composition of Benefits

Estimations of the demand for formal care confront an especially challenging simultaneity problem. As always, estimating the slope of demand requires an instrument that shifts supply but not demand. The key challenge in the case of formal care is that many factors that shift supply are also likely to shift demand by changing the opportunity cost of informal care. For example, consider using minimum wage laws (or their changes over time) as instruments for the price of formal care. Many formal home care workers earn roughly the minimum wage, so changes in the minimum wage are likely to shift the supply of formal care. But at the same time, changes in the minimum wage are also likely to change the opportunity cost of informal care-giving by changing the wage or employment prospects of some potential informal care-givers. Changes in the supply of informal care likely shift the demand for formal care since formal care and informal care are closely-related goods. The close links between formal care and informal care make it especially difficult to find valid instruments for the supply of formal care.

We circumvent this identification issue by taking advantage of the randomization in the Cash

and Counseling experiments. People randomized to the cash benefit face the market price of formal care, whereas people randomized to the traditional Medicaid program face a price of zero.¹³ Moreover, the price variation from the experiment is not only exogenous to the demand for formal care, it also spans the full range of prices most relevant for policy, from the market price down to zero.

The Cash and Counseling experiments were conducted by the Medicaid programs of Arkansas, Florida, and New Jersey. Participants were enrolled beginning in 1998 in Arkansas, in 1999 in New Jersey, and in 2000 in Florida. In New Jersey and Florida, only people who had been receiving Medicaid home care were eligible for the experiment, whereas in Arkansas anyone who was eligible for Medicaid home care could enroll. Although the exact implementations of the experiment differed somewhat between the three states, each shared the same basic features. After completing a baseline survey, participants were randomly assigned to one of two equal-sized groups. Members of the control group remained subject to the rules of the traditional, in-kind Medicaid home care program. Members of the treatment group were given cash budgets that could be spent on a wide range of personal care goods and services, including assistive devices, home modifications, and, most important, informal care from family or friends. Although the benefit had to be spent on personal care services, for many recipients it was likely to be equivalent to a pure cash transfer, since the vast majority of recipients had been receiving enough informal care at baseline to more than exhaust the budget they received from Medicaid.¹⁴

We use data from the nine-month follow-up survey for respondents who were at least 65 years old and who had non-missing data on their demographics (age, gender, race, and educational attainment) and self-rated health. Our sample contains 860 participants from Arkansas (428 in the cash group), 482 in Florida (245 in the cash group), and 604 in New Jersey (312 in the cash group). Table 1 provides summary statistics. At baseline, average formal care consumption ranged from 8 (Arkansas) to 16 (New Jersey) hours per week, and on average participants had two informal caregivers. The remaining rows of Table 1 show that the treatment and control groups are balanced on demographics such as gender, age, race, and education. Although non-negligible fractions of the treatment and control group

¹³ Appendix B.2 presents evidence that the marginal value of Medicaid formal care is zero for most recipients of the traditional in-kind benefit. It also tests the robustness of the estimated price sensitivity to alternative assumptions about the marginal value of Medicaid formal care.

¹⁴ Recipients had to submit receipts documenting that they spent their benefits on personal care services, though they were also allowed to spend up to 10 percent of their allowance on services that could not be readily invoiced (e.g., payments to a neighbor for mowing the lawn). The cash allowances received by members of the treatment group were slightly smaller than those that would have enabled them to purchase the care they would have received under the traditional Medicaid program. This is mainly because of a requirement that the experimental cash treatment be budget-neutral. This requirement meant that the costs of paying the counselors who helped treatment group members manage their care came out of the cash allowances.

attrited from the experiment before the nine-month follow-up survey (20 and 35 percent of treatment and control group members, respectively), the attrition did not lead to any apparent imbalance between the treatment and control groups in terms of observables. Of the 30 balance tests, none of the differences between treatment and control groups are statistically distinguishable from zero at the 5 percent level and only two are significant at the 10 percent level. These are fewer significant differences than would be expected to arise by chance without any differential attrition.

A simple measure of the sensitivity of the demand for formal care to the composition of benefits is the difference in average formal care consumption by people randomized into cash versus in-kind benefits. Table 3 makes this comparison. On average, people randomized to in-kind benefits consumed over twice as much formal care as people randomized to cash benefits: 14.8 versus 7.1 hours per week. Similarly large differences occurred in each state. Figure 2 shows the empirical cumulative distribution functions of formal care consumption by the in-kind and cash groups. The distributions have two key features. First, formal care consumption by the in-kind group is greater throughout the distribution. Second, many people consume no formal care, especially in the cash group. Whereas about one-fifth of the in-kind group consumes no formal care, over half of the cash group consumes no formal care. Receiving cash as opposed to in-kind benefits leads many people to a corner solution in which they consume no formal care. The prevalence of people consuming no care means that the observed mean differences between the two groups tend to understate the price sensitivity of demand. Overall, these simple comparisons suggest that the demand for formal care is quite sensitive to the composition of benefits.

The slope of the demand curve is not simply the ratio of the observed difference in average quantities to the difference in prices for two reasons. First, as just discussed, many individuals consume zero hours of care. We account for this by treating an individual's observed hours of care, q_i , as the outcome of a censored, latent demand for care, $q_i = \max\{0, q_i^*\}$. Second, some of the participants in the cash group reverted to traditional Medicaid home care, and some of the participants in both groups left Medicaid home care altogether. The differences in Table 3 therefore correspond to intent-to-treat parameters rather than direct measures of price sensitivity. We handle this issue by instrumenting for the price each participant faces with her randomized assignment.¹⁵ We estimate the system

$$q_i^* = \alpha + \beta p_i + X_i \gamma + \varepsilon_i$$

$$q_i = \max\{0, q_i^*\}$$

¹⁵For individuals who leave Medicaid home care or are in the cash group, the price they face is the market price in their state. For individuals who are in the traditional Medicaid home care program, the price they face is zero.

$$p_i = \mu_0 + \mu_1 \text{Cash}_i + X_i \mu_2 + \nu_i,$$

where p_i is the price of formal care, Cash is an indicator of whether the participant was randomized to the cash treatment, and X_i includes indicators for gender, education level, race, self-rated health, five-year age bins, and state. We begin by assuming (ε_i, ν_i) are jointly normal and estimate this system using an instrumental variables Tobit specification.

The first-stage results are presented in Table 4. On average, being assigned to the cash group increases the price of formal care by \$8.84, 64 percent of the average market price of \$13.73. The point estimate is fairly precise and the first-stage relationship is strong; the F-statistic exceeds 1,000, well above the levels at which weak instruments become a concern (Stock and Yogo, 2002). As one would expect given the random assignment, adding control variables has little effect on the estimated relationship between treatment assignment and the price of care (column (2)).

The instrumental variables estimate of β is presented in Table 5. The estimate implies that a one dollar increase in the hourly price of formal care reduces consumption by 1.8 hours per week. Evaluated at the sample means, this implies an elasticity near -1.2. When the demographic controls are added to the specification (column (2)), the estimate is virtually unchanged.

This price sensitivity suggests that formal care subsidies could potentially significantly increase consumption of formal care, as many people who otherwise without subsidies would have relied more heavily on informal care substitute toward formal care. This also suggests that subsidizing formal care may have large moral hazard costs. For someone consuming the average amount of formal care among people randomized to traditional in-kind benefits (14 hours per week), the ex-post deadweight loss from in-kind provision is about 72 percent of Medicaid’s spending on the individual’s care.¹⁶ For someone consuming 25 hours of formal care per week, the deadweight loss from in-kind provision is about 49 percent of Medicaid’s spending on the individual’s care. More generally, and as discussed in Section 2, the ex-post deadweight loss as a fraction of Medicaid spending is decreasing in the level of demand for care.

In Appendix B.2, we analyze the consequences for the estimated price sensitivity of making a wide range of alternative assumptions about the distribution of the error terms and limits on how much Medicaid home care recipients of the traditional in-kind benefit can consume. The conclusion that the demand for formal care is quite sensitive to its price is robust to a wide range of alternative assumptions.

¹⁶With a price sensitivity of -1.8, someone consuming 14 hours of care per week has an equivalent variation of formal care benefits (assuming no income effects) of \$54 per week, or 28 percent of Medicaid’s \$192 of spending on that care (14 hours per week at an average price of \$13.73 per hour).

Also in Appendix B.2, we discuss the generalizability of the estimates of the demand for formal care based on the Cash and Counseling experiments to other policies and populations of interest. There are two key issues. First, selection into the experiment likely means that the average participant in Cash and Counseling was more price sensitive than the average recipient of Medicaid home care and potentially either more or less price sensitive than the average person in the population of people eligible for Medicaid home care. Second, the nature of the experiment—especially its unexpected nature and uncertain duration—seems likely to lead to smaller quantity responses than one would expect from permanent changes in policies. Although these are important issues for generalizability, the benefits of the large, exogenous variation in the price of formal care make the Cash and Counseling experiments an important source of evidence on the demand for formal care. In light of the possible issues with generalizability, however, in the policy analysis (Section 5) we test the robustness of our results to a wide range of values of the price sensitivity.

5 Heterogeneity in the Demand for Formal Care

In this section, we estimate the heterogeneity in the demand for formal care and assess the extent to which it is correlated with observables that could potentially serve as tags.¹⁷ We do this using the 1999 wave of the National Long Term Care Survey (NLTC). The NLTC is a nationally representative sample of Medicare enrollees, including people in care-giving institutions such as nursing homes. It is designed to provide information on the health and functional status of Americans who are at least 65 years old. We focus on people with at least two activities of daily living limitations who are living in the community because this is the population of interest for home care benefits programs. The data include information on formal care consumption, income, physical limitations, and a host of demographic characteristics.

Heterogeneity in formal care consumption implies, holding other resources fixed, heterogeneity in non-care consumption, which in many models implies heterogeneity in the marginal utility of wealth. Figure 3 shows the empirical cumulative distribution function of formal care hours. There are two important features of this distribution. First, many people do not consume any formal care—approximately 62 percent of the sample. Second, there is a long right tail of people who consume large amounts of care. Conditional on consuming any care, median consumption is 14 hours of care per week, whereas the 95th percentile is 168 hours

¹⁷In principle, even a characteristic that is uncorrelated with formal care consumption could be a valuable tag. Any characteristic that is predictive of higher moments of the distribution of formal care would be useful. But in practice, characteristics that are highly correlated with formal care are likely to be the most valuable tags.

of care per week.

To the extent that an insurance program can condition transfers on verifiable characteristics, the observed heterogeneity in Figure 3 overstates the relevant heterogeneity for the insurance program. For instance, many Medicaid home care programs require that an individual have at least two activities of daily living limitations to be eligible. Figure 4 plots the CDFs of formal care for people with fewer than two activities of daily living limitations and for people with at least two limitations. There is a clear difference in care use across the two groups; the number of activities of daily living limitations explains some of the variation in care use. But a large amount of heterogeneity remains even within these groups.

We turn to assessing the extent to which the heterogeneity in formal care can be explained by observable characteristics. We do so by regressing measures of formal care consumption on a set of observable characteristics that could be correlated with the demand for care. First, we run a probit regression of whether someone uses formal care as a function of age, indicators for her number of activities of daily living limitations (omitted group is two limitations), indicators for her self-rated health (omitted category is excellent), gender, whether she lives alone, whether she has children, and linear and quadratic terms for her income. The marginal effects from this regression are presented in column (1) of Table 6. Age, activities of daily living limitations, self-rated health, and living alone are all strongly associated with whether someone consumes formal care. In the second column, we present results of an OLS regression of hours of care per week on observables for the set of people who consume positive amounts of care. Physical limitations, gender, and whether the person lives alone are important predictors of care use among those consuming some formal care. In the third column, we use our estimates of the price sensitivity of demand for formal care to “correct” for differences in the prices people face and find similar results.¹⁸ The next three columns display results where state fixed effects have been included in each specification. The results closely mirror those without the state fixed effects.

For understanding the targeting properties of potential tags, it is important to understand not only the extent to which observable characteristics predict differences in mean consumption but also how much of the variation they explain. Table 7 reports McFadden’s pseudo R-squared statistics for the probits and standard R-squared statistics for the ordinary least squares regressions of formal care use on observed characteristics. The columns correspond to the same sets of analyses presented in Table 6, but the rows now indicate which variables have been added to the set of controls. The set of variables is cumulative so each row includes all of the previous rows’ controls as well.¹⁹ The observables predict a moder-

¹⁸Specifically, we subtract 1.8 times the market price for care from each individual’s observed care use except for those who face a price of zero, i.e. those on Medicaid.

¹⁹“Health controls” variables are age, indicators for the number of activities of daily living limitations,

ate amount of the variation in care use—between 7 and 28 percent—but leave most of the variation “unexplained.”²⁰

Taken together, these analyses suggest two main conclusions. First, there is a large amount of heterogeneity in the demand for care that cannot be explained by observable characteristics in the NLTCs. This implies that even if it were possible to condition insurance benefits on all of these characteristics at reasonable administrative and moral hazard costs, there would remain important residual heterogeneity that could potentially be targeted with in-kind formal care benefits. Second, certain observable characteristics, especially the number of activities of daily living limitations and whether the individual lives alone, are associated with large differences in care use on average. Such characteristics could potentially serve as tags to target benefits to desired recipients.

To get a rough sense of the potential scope for targeting benefits, we estimate the fraction of people eligible for Medicaid home care who actually take up benefits. As discussed in Section 2, in-kind transfers differentially transfer resources to high-demand types even when take-up is 100 percent. But to the extent that taking up benefits is costly, some targeting may occur on the extensive margin, as people with relatively low demand may not choose to join the program. Take-up rates are therefore an important metric of the extent to which providing benefits in kind targets transfers to high-demand types. Among people eligible for Medicaid home care, we estimate that between 4 percent and 16 percent actually take up benefits.²¹ The low take-up rate means that, holding fixed total spending on the program, benefits per recipient are between 6 and 24 times greater than they would be under a hypothetical program with 100 percent take up. Applying these ratios to participants in the Cash and Counseling experiments, a hypothetical program with 100 percent take up would transfer

indicators for self-rated health categories, and gender. “Informal care options” variables include whether the individual lives alone and whether she has children.

²⁰The variance in formal care consumption understates the variance in the demand for formal care because consumption must be non-negative; variance in consumption does not capture heterogeneity among people who do not consume any care. As a result, this analysis tends to overstate the fraction of the variance in the demand for formal care “explained by” these observables.

²¹Estimating take-up rates for Medicaid home care, and Medicaid more generally, is notoriously difficult (U.S. Department of Health and Human Services, 1992). Eligibility rules are complex, vary from state-to-state, and often depend upon household characteristics that are unobservable to the researcher. We use the NLTCs to estimate the fraction of the elderly who are eligible for benefits, based on the eligibility criteria from Schneider et al. (1999). We combine that estimate with the size of the sixty-five and older population and administrative estimates of the number of Medicaid home care users from LeBlanc et al. (2001). The main source of uncertainty in our estimated take-up rate (from four to 16 percent) is the incompleteness of the information on household assets in the NLTCs. In all cases, a person must have at least two activities of daily living limitations. The upper end of our range, 16 percent take-up, should be interpreted as an upper bound on the take-up rate because we imposed (much) more restrictive income and asset requirements than the actual limits in the vast majority of states. In particular, we imposed that the household earned no more than 100 percent of the SSI benefit and had no cars (car value is one of the primary inputs to the asset tests). Our least restrictive eligibility threshold uses the income limits from Schneider et al. (1999) and imposes that the household have fewer than two cars.

to each recipient the dollar value of just 0.6 to 2.3 hours of care per week rather than the 14 hours of care per week recipients actually received. This reinforces the conclusion that providing home care benefits in kind may have important targeting effects.

Whether this targeting of benefits to a subset of the eligible population is desirable depends on the nature of selection into the program and the planner's objective function. Table 8 shows summary statistics on the characteristics of the population that is eligible for Medicaid home care broken into those who do and those who do not actually take up the in-kind benefits. Individuals who take up Medicaid home care have significantly greater physical disabilities, less education, and are less likely to have a spouse who can provide informal care. Even after conditioning on the observable indicators of wealth and health status used by Medicaid to determine eligibility, there remains substantial variation in functional limitations and availability of informal care options, and people who actually take up Medicaid home care benefits appear to be in worse health and have less good informal care options than eligibles who do not take up benefits.

The significant heterogeneity in formal care consumption among people with the same observable characteristics, the low take-up rate of traditional Medicaid home care, and the observable differences between those who take up traditional Medicaid home care and the broader population of people eligible for the program all suggest that providing home care benefits in kind may have important targeting benefits.

6 Policy Analysis

In this section we adapt an approach from the literature on optimal taxation to evaluate the welfare consequences of alternative home care benefit programs. We consider programs that combine a formal care subsidy with a cash benefit component. As discussed in Section 2, the key determinants of the welfare consequences of in-kind provision of formal care are the distribution of demand for care, the sensitivity of the demand for care to the composition of benefits, and the link between marginal utility and the demand for care. The key source of heterogeneity is in the demand for formal care, which may be driven by hard-to-verify heterogeneity in health status, informal care options, or abilities to cope with health problems. We also test the robustness of the results to different assumptions about how unverifiable heterogeneity in utility might be correlated with demand for formal care (state-dependent utility).

6.1 Model and parameterization

Individuals draw their type from a distribution of types, $\theta \sim F(\theta)$. Then they choose their formal care consumption and non-care consumption to maximize utility subject to a budget constraint that depends on the policy in operation. The budget constraint is

$$A + pF = m,$$

where F is formal care consumption, A is non-care consumption (i.e., “all other goods,” the numeraire), p is the after-subsidy price of formal care, and m is benefit-inclusive nominal income. The utility function and the corresponding demand for care are

$$U(A, F; \theta) = u \left(A - \frac{(\max\{\alpha, 0\} - F)^2}{2\beta} \right),$$

$$F(p, m; \theta) = \max \{0, \min \{m/p, \alpha - \beta p\}\}.$$

α is the quantity of care at which the individual is satiated, i.e., the amount of care the individual would consume when facing a price of zero. β determines the utility cost of consuming levels of care other than the satiation level and thereby determines the sensitivity of the demand for formal care to the composition of benefits. $F(p, m; \theta)$ is the Marshallian demand function for formal care.

This utility function is motivated by key evidence from our setting. It produces a simple function for the demand for formal care that is consistent with the sensitivity of formal care consumption to its price and that people become satiated at finite levels of formal care consumption.²² This utility function also has several appealing features. It nests as a special case the widely-used model in which health spending is equivalent to a wealth shock.²³ It implies that the demand for formal care is linear in its price within the range of prices in which the individual is not at a corner. It has an intuitive interpretation: Utility is decreasing in any unmet, residual health needs, $(\alpha - F)$, the size of which is decreasing

²²The most direct evidence of satiation is that among Cash and Counseling participants with information on their care plan hours, 43 percent consumed less care than they were entitled to based on their care plan. Intuitively, satiation might arise from a demand for privacy or space, since home care involves close contact with caregivers in one’s home. This utility function is also consistent with the fact that most people who need assistance do not consume any formal care. This implies that there is no Inada condition on formal care consumption and that formal care is not too complementary with other goods that people consume.

²³As β approaches 0, formal care consumption approaches α ($F(p, m; \theta) \rightarrow \alpha$, ignoring corner solutions), and the indirect utility function approaches $v(p, m; \theta) = u(m - p\alpha)$. For $\beta > 0$, demand for formal care is sensitive to its price and the indirect utility function is $v(p, m; \theta) = \begin{cases} u \left(m - \frac{\max\{\alpha, 0\}^2}{2\beta} \right), & \text{if } \alpha < \beta p; \\ u \left(m - p(\alpha - \beta p) - \frac{\beta p^2}{2} \right), & \text{if } \alpha \geq \beta p. \end{cases}$

This differs from the benchmark case in which health spending is a wealth shock by just a slight adjustment, which is necessary to accommodate the observed price sensitivity of demand for formal care.

in formal care consumption, F , and increasing in the level of demand for formal care, α . This captures the idea that certain health problems are costly for people to cope with on their own. Such costs create a demand for formal care, which reduces the utility costs of bad health. The greater is the price of formal care, the less of it that people will consume, and so the greater will be the residual utility costs of coping with health problems left over after formal care consumption. Marginal utility of income depends on the demand for formal care mainly through the budget constraint: Greater spending on formal care means lower non-care consumption. The key assumption is that formal care is not a good substitute for non-care consumption as a whole.

To better understand the nature of the risk that people face and the desired insurance transfers, consider the benchmark of a first-best insurance program. The first-best transfer schedule satisfies:

$$b(\theta; B) = \begin{cases} b(B) + \frac{\max\{\alpha, 0\}^2}{2\beta}, & \text{if } \alpha < \beta p; \\ b(B) + p(\alpha - \beta p) + \frac{\beta p^2}{2}, & \text{if } \alpha \geq \beta p, \end{cases}$$

where B is average per-person spending on people eligible for home care benefits and $b(B)$ is the cash transfer that makes total program spending equal B . The first-best transfer is increasing in α , first quadratically then linearly. With these transfers, indirect utility is

$$v_{FB}(p, m, B; \theta) = u(m + b(B)),$$

which is independent of θ . The first-best contract does not distort consumption, and it fully insures all risk. By making larger transfers to people with larger demands for formal care, it fully compensates people for their expenditures on formal care and any residual utility costs they face from coping with their health problems.

6.2 Parameter values

Everyone has the same price sensitivity of demand for care, β . The only heterogeneity is in the level of demand for care, α . For the price sensitivity of demand for formal care, we use our main estimate from the Cash and Counseling experiment, $\hat{\beta} = 1.8$. This estimate implies that every \$1 increase in the hourly price of formal care reduces formal care consumption by 1.8 hours per week. We test the robustness of the results to a wide range of alternative values.

We use $\hat{\beta}$ to convert the joint distribution of formal care consumption and formal care prices observed in the NLTCs into a distribution of the level of demand for formal care, $F(\alpha)$.

For the analysis that takes as given standard eligibility criteria for home care benefits, our sample is everyone aged 65 and older with at least two activities of daily living limitations. For the tags analysis, we estimate separate $F(\alpha)$ distributions for each sub-group of this population as defined by their tagged characteristics, for example, separately for people with different numbers of activities of daily living limitations. Estimating $F(\alpha)$ would be entirely straightforward were it not for people who consume no care when facing a positive price—62 percent of the population of interest. Among people who consume no formal care, revealed-preference analysis implies only a bound on the level of their demand (that their marginal value at zero hours of care is no greater than the price). But because we will be analyzing policies that reduce the prices people face, it is important to know at which price each individual would begin purchasing care. We handle this fundamental unobservability issue by extrapolating the observed distribution among people who consume a strictly positive amount of care backward to “fill in” the unobservable α values of people who consume no formal care when facing a positive price. Details of this calculation are reported in Appendix C. We also test the robustness of the results to a wide range of alternative assumptions about this portion of the α distribution.

Figure 5 presents our main estimate of the distribution of the level of demand for formal care, $F(\alpha)$. The key features of this distribution, inherited from the observed distribution of formal care consumption, are that it exhibits a long right tail (the mean far exceeds the median) and that most of the mass is at low values. The estimates imply that 0 percent of people would consume no care even at a price of zero, whereas the observed distribution of formal care when people face market prices has 62 percent of its mass at zero hours of care.²⁴ The difference comes from the estimated price sensitivity of demand for formal care together with the shape of the observed formal care distribution among people who consume a strictly positive amount of care when facing market prices.

We adopt standard values of the remaining parameters. 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; De Nardi et al., 2010; Ameriks et al., 2011). In our model, the argument c is “net consumption,” non-care consumption net of any residual coping costs, $c = A - \frac{(\alpha - F)^2}{2\beta}$. We follow Brown and Finkelstein (2008) in taking as a baseline value a coefficient of relative risk aversion, γ , of three. Income before transfers, m , is \$15,000

²⁴Someone with $\alpha_i < 0$ would have to be paid in order to consume formal care. The difference between $\alpha_i = 0$ and $\alpha_i < 0$ is only relevant for formal care subsidies in excess of 100 percent. Although it might at first appear unintuitive that some people would consume no care when facing a zero price, it is important to keep in mind the nature of home care. Home care is a specific set of services that are often private and sensitive, such as assistance with bathing and using the toilet. Much of it is not just general help around the home that even healthy people would value. An important caveat to this analysis is that as formal care subsidies increase, the incentive to relabel other services as formal care gets larger. Large enough subsidies would presumably have to be accompanied by enforceable limits on the set of services that qualify.

per year. The distribution of before-subsidy prices of formal care is the empirical distribution observed in the NLTCs. People who cannot achieve net consumption of at least $\bar{c} = \$5,000$ per year receive transfers that enable them to enjoy net consumption of \$5,000 per year (a consumption floor).

6.3 Optimal formal care subsidy given standard eligibility criteria

In this section we calculate the optimal subsidy rate on formal care, σ^* , taking as given total spending on program benefits and standard eligibility criteria for home care benefits.²⁵ We focus on programs that limit eligibility to people with two or more activities of daily living limitations. This follows standard practice for Medicaid home care and private long-term care insurance.²⁶ Everyone who qualifies for home care benefits automatically receives the same cash benefit and the same subsidy rate on their purchases of formal care. These choices focus our analysis on the core tradeoff at the heart of in-kind provision. We consider policies under which total program spending equals the spending on a pure in-kind benefit program, a 100 percent subsidy with no cash benefit. Policies with smaller subsidy rates have larger cash benefits.

Figure 6 presents the key result. It shows the equivalent variation of the mixed in-kind and cash benefit policy as a function of the in-kind component, the subsidy rate σ . The equivalent variation is expressed as the pure cash transfer (identical for everyone who qualifies for home care benefits) that would be required to make the individual as well off as she is under the mixed-benefit program in question. The optimal subsidy rate is 88 percent, close to a pure in-kind program (under which the after-subsidy price is zero). The optimal subsidy increases welfare substantially relative to a pure-cash benefit program. In order to make the individual as well off as she is with the optimal policy under an alternative pure-cash benefit program, the cash benefit would have to be about 80 percent greater than the average cost of the in-kind program. Figure 6 also shows, however, that the optimal subsidy is significantly less valuable than the hypothetical first-best policy full insurance. The optimal in-kind subsidy achieves 59 percent of the incremental value over a pure-cash benefit that the first-best policy does. Table 7 shows a variety of outcomes for each of several different versions of the model. The

²⁵The full optimal policy problem also involves choosing the eligibility criteria on which to condition benefits (e.g., the program might offer different benefit bundles to people with different numbers of activities of daily living limitations) as well as the net transfers across verifiably-distinguishable groups. We address these issues in the next section. A socially- (as opposed to just privately-)optimal program would also take into account any impacts of the program on real and fiscal externalities, including any effects on the spending of consumption-floor programs and on tax revenue.

²⁶In analyzing formal care subsidies for which some people are not eligible, the impossibility of reselling formal care benefits is important. This is an important way in which formal care differs from food stamps, for example.

purpose of this table is to provide intuition for and assess the robustness of the key results. The first column of the table shows results for the baseline specification just discussed. The key tradeoff involved in increasing the in-kind component of the benefit can be seen clearly by comparing the average level of and dispersion in non-care consumption under the optimal subsidy program and under the cost-equivalent pure-cash benefit program. The optimal subsidy reduces average non-care consumption due to the consumption distortion (and to a lesser extent due to foregone transfers from the consumption floor), but it also greatly reduces the dispersion of non-care consumption, as measured by the standard deviation. Under the pure-cash program the standard deviation of annual non-care consumption is 4.5 times greater than under the optimal program, \$5610 vs. \$1237.

Additional rows of the table unpack these results further. They show that formal care consumption is significantly greater under the optimal subsidy than in the absence of any program, by a factor of 2.4. This translates into a large distortion cost; the total ex-post equivalent variation of the optimal program summed over all recipients is only 48 percent of the total cost of the program. Part of this is due to the optimal program displacing transfers from the consumption floor, but much of it is due to the consumption distortion from the formal care subsidy. This amounts to a significant implicit tax on insurance, equivalent to a tax of almost 100 percent of benefits. The reason that subsidizing formal care is optimal despite the large distortion is that the in-kind subsidy redistributes toward people with greater marginal utility. The correlation between an individual's marginal utility in the absence of any program and his ex-post equivalent variation of benefits under the optimal program is 0.84. The net benefit from in-kind provision comes from making large transfers to the relatively few people with high demand for care (and so low non-care consumption). This can be seen in the bottom row of the table, which shows that under the optimal subsidy program only 16 percent of people value the subsidy program as much as the cost-equivalent pure-cash benefit program ex post. This may help explain why many countries and U.S. states have made home care benefits more cash-like. Making benefits more cash-like helps most individuals ex post, often significantly. A key finding of this paper, however, is that the greater ex-post value of more cash-like benefits comes at the expense of much less redistribution toward people with high demand for formal care, which may worsen insurance.

The other columns of the table test the robustness of the results to making different assumptions about the key ingredients of the model. The price sensitivity of demand for formal care must be quite large—over 10 times larger than we estimate based on evidence from the Cash and Counseling experiment—in order to overturn the conclusion that the optimal subsidy is large. Even if the distribution of partially-identified α values among people who consume no care at positive prices is in the “worst-case” configuration (i.e., each α_i equal to the maximum value consistent with observed behavior), the optimal subsidy rate is still

86 percent.²⁷ The utility function must exhibit strong state dependence of just the right kind—greatly decreasing the marginal utility of people with high demand for formal care in just the right way²⁸—in order to overcome the fact that, holding other resources constant, people with greater formal care consumption must have lower non-care consumption. Although the right tail of the distribution of demand for formal care is an important determinant of the targeting benefit and so the optimal subsidy, the optimal subsidy remains large even when the right tail of the distribution is chopped off or when all of the α values are scaled down. If people who consume more than 50 hours per week of care are dropped from the simulation, the optimal subsidy is 59 percent. If all of the α values are cut in half, the optimal subsidy is 75 percent.²⁹ Finally, a combination of relatively low risk aversion together with a relatively generous consumption floor can overturn the optimality of a large subsidy on formal care, although this reflects the undesirability of *any* insurance—including a first-best contract—in situations in which means-tested programs are sufficiently attractive rather than any undesirability of in-kind benefits per se.³⁰

The alternative specifications, in addition to providing a sense of the robustness of the

²⁷The intuition for the robustness of the results to the exact shape of the distribution of demand among people with a low demand for formal care is that the key driver of the targeting benefit from in-kind provision is the shape of the other tail of the formal care distribution: people with high demand for care. The distribution of demand among people with a low demand for care matters mainly for determining the distortion cost of in-kind provision.

²⁸Appendix C discusses how different types of state-dependence of utility affect the desirability of providing home care benefits in kind. That many private long-term care insurance contracts subsidize the consumption of formal care suggests that the correlation between marginal utility without insurance and formal care consumption is positive, at least for many of the people who buy these contracts. This suggests that to the extent that there is state-dependence of utility, it does not overturn the positive correlation between formal care consumption and marginal utility that would tend to arise from the budget constraint alone (since greater consumption of formal care reduces the resources available for non-care consumption).

²⁹This partially addresses possible biases from modeling a dynamic situation in a static model. The static nature of the model means that 100 percent of formal care costs must be financed by reducing non-care consumption in that period; formal care costs cannot be smoothed over time by saving and borrowing. To the extent that shocks are not entirely persistent, this tends to lead us to overstate the welfare cost of uninsured risk and so the value of insurance against it. This issue is less relevant for Medicaid home care—with its strict asset tests—than for private long-term care insurance. It also addresses possible biases from ignoring other risk-sharing arrangements, e.g., informal family insurance.

³⁰The final column shows that if risk aversion is relatively low ($\gamma = 1$) and the consumption floor is relatively generous ($\bar{c} = \$5,000$), the first-best insurance policy that provides complete insurance without distorting consumption is dominated by an alternative uniform pure-cash benefit that provides no insurance at all. The reason that even a first-best, actuarially-fair insurance contract is dominated by the no-insurance alternative in this case is the high rates of implicit taxation from the consumption floor. Without insurance, the consumption floor pays for much of the care of people with the greatest demand for care. As a result, insurance reduces average consumption among the insured by reducing the transfers they receive from consumption-floor programs. This is similar to Brown and Finkelstein's (2008) findings about how Medicaid can crowd out purchases of even actuarially fair long-term care insurance by a large part of the wealth distribution. It should be noted that while the first-best contract is dominated by no insurance from the perspective of people eligible (or potentially eligible) for home care, the first-best contract is better from the perspective of society as a whole. From the perspective of society as a whole, the home care benefit should internalize any effects alternative home care benefits might have on the rest of society, including government or private consumption-floor programs.

results, also provide information about the key factors driving the results. As expected, the net benefit of subsidizing formal care is decreasing in the price sensitivity of demand for formal care. When demand for formal care is completely inelastic ($\beta = 0$), a 100 percent subsidy achieves the first best.³¹ The targeting benefit of in-kind provision is increasing in risk aversion and decreasing in the generosity of alternative insurance arrangements, such as any consumption floor or means-tested programs. The targeting benefit of subsidizing formal care is increasing in the extent to which there is state-dependent utility in which marginal utility is greater for people with greater demand for formal care (above and beyond the effects operating through the budget constraint or residual coping costs). If such state-dependence is strong enough, it is optimal to more than fully subsidize formal care (columns 8 and 11).

Although formal care subsidies significantly increase risk sharing, they (optimally) leave some risk uninsured due to the distortion they cause. Both the incompleteness of the insurance and the distortion from the subsidy mean that formal care subsidies fall short of achieving the first best. In the baseline specification, the optimal subsidy achieves about 83 percent of the incremental value over pure-cash benefits of the first-best policy. The shortfall is a measure of the potential gain from using a richer set of policies. A natural enrichment is to condition benefits on verifiable characteristics—i.e., to use tags,—a possibility to which we now turn.

6.4 Optimal policy with tags

This section extends the analysis to the case in which different groups of people, defined by their verifiable characteristics, can be offered different benefits. We estimate the gains from catering benefits to different groups of people defined by the number of activities of daily living limitations they have (2–4, 5, and 6) and whether they live alone, the two strongest predictors of formal care consumption uncovered in Section 5.³² The procedure is the same as that in the last section, except that we estimate different α distributions for different groups of people and allow the program to offer different benefits to people in different groups. Figures of the α distributions of each group are reported in Appendix C.

³¹One caveat about this result is that it is based on a model in which formal care is borderline inferior (no income effects). This result need not hold in a more general model with income effects of demand for formal care. It is also important to note that the assumption that formal care is borderline inferior tends to work against the value of in-kind provision by increasing the consumption distortion. The greater are income effects of demand for formal care, the more that the (negative) income effects from subsidizing formal care (due to the consumption distortion) offset the inefficient over-consumption of formal care due to the substitution effect.

³²We are limited in the number of groups we can split the population into by the size of the NLTCs sample. We chose the groups to maximize the across-group heterogeneity in the demand for formal care.

Table 7 shows the results. The planner takes advantage of the flexibility offered by tagging to redistribute across groups by offering high-demand groups more generous benefits. In the case of a tag based on whether someone lives alone, for example, the planner transfers on average over two-and-a-half times more to people who live alone than to people who live with others. Despite this, however, the ex-ante welfare gain from using these tags is quite small. The incremental welfare gain from optimally tagging a pure-cash benefit based on whether someone lives alone is \$227, just 4 percent of the gain from an optimal untagged mixed benefit. The incremental welfare gain from optimally tagging a mixed cash-in-kind benefit based on whether someone lives alone is even smaller—just \$44. The incremental gains from tagging benefits based on the number of activities of daily living limitations someone has are smaller still. The fundamental reason for tags’ ineffectiveness in insuring this risk is that much of the heterogeneity in demand for formal care occurs within rather than across groups that can be distinguished on the basis of their verifiable characteristics. This comes through in the correlation between someone’s marginal utility in the absence of any program and the optimal tagged pure-cash benefits. This correlation is just 0.2 when whether someone lives alone is used as a tag and just 0.05 when someone’s number of activities of daily living limitations is used as a tag.

These results suggest that in the context of homecare, the potential for tags to target transfers toward high-marginal utility types is limited. These results are consistent with those of Mankiw and Weinzierl (2010) on the effects of using height as a tag for optimal income taxation.³³ Although different combinations of observable characteristics could potentially improve on those we have analyzed here, both the small gains from tags based on two of the strongest predictors of formal care consumption and the limited extent to which the full set of observable characteristics in the NLTCs predict formal care consumption (as discussed in Section 5) suggest that the scope for tags is limited and reinforce the conclusion that in-kind benefits have an important role to play in terms of targeting benefits to high-marginal utility types.

6.5 Summary of results

This analysis is subject to several caveats. It assumes that people’s decisions about consumption are rational. This ignores possible paternalistic rationales for in-kind transfers, which could be important in the case of home care given the cognitive health problems from which some recipients suffer. Such considerations, which were one of the main motivations

³³In both cases, the optimal tagged transfers are large; the optimal “lives-alone subsidy” is \$4790 and the optimal “height tax” on someone earning \$50,000 is \$4,500. But the welfare gains from tagging are a small fraction of aggregate income—about 1.5 percent for a “lives-alone subsidy” and about 0.2 percent for a “height tax.”

for the Cash and Counseling experiments, seem likely to increase the value of in-kind as opposed to cash transfers in this context. It assumes that take-up decisions are fully rational and abstracts from any costs of taking up. This is done to focus on the core tradeoff at the heart of in-kind provision, but it is important to note that many in-kind programs have low take-up rates, whether from low knowledge about or high costs of taking up the programs. It assumes that all ex-post heterogeneity is the outcome of an exogenous process. This feature, which is shared by the vast majority of the large literature on optimal taxation, rules out ex-ante moral hazard (effects of policies on the distribution of ex-post types), which tends to increase the net value of insurance or redistribution. It focuses only on home care and does not explicitly model substitution across other types of care. This was done for simplicity given that there appears to be little substitution across different types of long-term care (Grabowski and Gruber, 2007; Kemper, 1988). Finally, it focuses on singles in order to avoid the many complexities involved in modeling couples, including any financial risk sharing and utility consequences of different caregiving arrangements. An analysis of couples and extended families is an interesting topic for future work.

Acknowledging these caveats, taken as a whole the results suggest that in-kind provision of formal care benefits likely increases welfare despite the large distortion it causes. Although the distortion cost of this particular means of targeting is large, our results suggest that the main alternative means of targeting (using tags) is unlikely to be very effective in this context. These conclusions are robust to a wide range of assumptions. The fundamental reasons for this robustness are the large extent of hard-to-verify heterogeneity in the demand for formal care and the rapid rate at which marginal utility diminishes in the level of consumption under standard utility functions (Kaplow, 2011).

7 Conclusion

We analyze the consequences of providing formal home care benefits in kind, focusing on the targeting benefits and distortion costs. In-kind provision appears to increase welfare despite imposing large distortion costs, since the targeting benefits are even larger. The key factor driving this result is the significant, hard-to-verify heterogeneity in the demand for formal care—whether from hard-to-verify differences in underlying health or in the costs of coping with a given set of health problems—which implies significant heterogeneity in non-care consumption and so, in many models, in marginal utility.

Two main caveats are important to keep in mind in interpreting our results. First, the magnitude of the targeting benefit depends crucially on the utility function, particularly any (state-)dependence of utility on the unverifiable heterogeneity that motivates in-kind

provision in the first place. This caveat, common to all questions about the optimal design of insurance programs and other policies that redistribute across people, is extremely important: Improving our understanding of the mapping between spending on market goods and utility is a high priority for future work. Second, we focus only on the potential targeting benefit of providing formal care in kind. We do not consider other potential benefits, including improving tax system efficiency, alleviating the Samaritan's dilemma, and paternalistic benefits, all of which might be important in this context. These other potential benefits seem likely to increase the net value of in-kind provision of formal care, which would reinforce our conclusion about the desirability of in-kind provision of formal care.

Our results have important implications for understanding long-term care risk, private long-term care insurance markets, and public long-term care insurance programs. Long-term care risk, at least the portion of the risk related to home care, appears to feature a combination of important hard-to-verify heterogeneity and high price sensitivity, which makes it an especially difficult risk to insure. Our results imply that even optimally-designed home care insurance involves a significant distortion cost, which acts as a large implicit "moral hazard" tax on insurance. This implicit tax, together with numerous other factors that have been identified in the literature (see Brown and Finkelstein, 2011, for a review), could help explain the small size of markets for private long-term care insurance.

Our results also have important implications for public policies, especially recent policy initiatives to make various in-kind benefits more flexible and cash-like. The Medicaid Cash and Counseling programs and the short-lived CLASS Act component of what went on to become the Affordable Care Act are just two examples among many of "consumer-driven" reform proposals that aim to increase the range of options available to benefit recipients. The main impetus for these proposals is the view that the consumption distortion from in-kind provision is large, a view that is consistent with our analysis of the particular case of Medicaid home care benefits. Yet a frequently-overlooked consequence of such reform proposals is that, in addition to reducing consumption distortions, such proposals would also tend to systematically change the distribution of benefits received by different people. Our analysis suggests that although moving toward more cash-like home care benefits would significantly reduce consumption distortions, it would also forgo the apparently even larger targeting benefits that arise from providing formal care in kind. To the extent that achieving a good targeting of benefits in any particular context is difficult or infeasible without distorting consumption, as our analysis suggests may be the case in the context of home care, any gain from reducing distortions must be weighed against any reduction in targeting efficiency that would result.

The issue of optimal benefit design in government programs is a central one, as many of the

most important government programs involve in-kind benefits, including public schooling, food stamps, public housing, and Medicare and Medicaid medical benefits. Although home care shares much in common with other types of health care, the quantitative magnitudes of the key factors determining the desirability of in-kind provision that we emphasize—the distribution of demand and the sensitivity of demand to benefit composition—will be different in different contexts. It is therefore important to evaluate the costs and benefits of alternative benefit designs on a case-by-case basis, and our hope is that the approach we have developed in this paper will prove fruitful in the analysis of other policies as well.

References

- Akerlof, G. A. (1978). The economics of “tagging” as applied to the optimal income tax, welfare programs, and manpower planning. *The American Economic Review* 68(1), 8–19.
- Ameriks, J., A. Caplin, S. Laufer, and S. V. Nieuwerburgh (2011, 04). The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives. *Journal of Finance* 66(2), 519–561.
- Arno, P. S., C. Levine, and M. M. Memmott (1999). The economic value of informal caregiving. *Health Affairs* 18(2), 182–188.
- Barczyk, D. and M. Kredler (2016). Evaluating long-term care policy options, taking the family seriously. *Review of Economic Studies*, forthcoming.
- Blackorby, C. and D. Donaldson (1988). Cash versus kind, self-selection, and efficient transfers. *The American Economic Review*, 691–700.
- Brown, J. R. and A. Finkelstein (2007, November). Why is the market for long-term care insurance so small? *Journal of Public Economics* 91(10), 1967–1991.
- Brown, J. R. and A. Finkelstein (2008). The interaction of public and private insurance: Medicaid and the long-term care insurance market. *The American Economic Review* 98(3), 1083–1102.
- Brown, J. R. and A. Finkelstein (2011). Insuring long-term care in the United States. *Journal of Economic Perspectives* 25(4), 119–42.
- Brown, R. S., B. L. Carlson, S. Dale, L. Foster, B. Phillips, and J. Schore (2007). *Cash and Counseling: Improving the Lives of Medicaid Beneficiaries who Need Personal Care Or Home-and Community-based Services: Final Report*. Mathematica Policy Research.
- Bruce, N. and M. Waldman (1991). Transfers in kind: Why they can be efficient and nonpaternalistic. *The American Economic Review* 81(5), 1345–1351.
- Centers for Medicare and Medicaid Services (2010). National health expenditures by type of service and source of funds, cy 19602008. Technical report.
- Cunha, J. M., G. De Giorgi, and S. Jayachandran (2011). The price effects of cash versus in-kind transfers. Technical report, National Bureau of Economic Research.
- Currie, J. and F. Gahvari (2008, June). Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature* 46(2), 333–383.
- Da Roit, B. and B. 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.
- De Nardi, M., E. French, and J. B. Jones (2010, 02). Why do the elderly save? The role of medical expenses. *Journal of Political Economy* 118(1), 39–75.
- Doty, P., K. J. Mahoney, and M. Sciegaj (2010). New state strategies to meet long-term care needs. *Health Affairs* 29(1), 49–56.

- Ettner, S. (1995, February). The impact of “parent care” on female labor supply decisions. *Demography* 32(1), 63–80.
- Foster, L., R. Brown, B. Phillips, J. Schore, and B. Lepidus Carlson (2003, March). Improving the quality of medicaid personal assistance through consumer direction.
- Friedberg, L., W. Hao, W. Sun, A. Webb, and Z. Li (2014, October). New evidence on the risk of requiring long-term care. Working Papers, Center for Retirement Research at Boston College wp2014-12, Center for Retirement Research.
- Grabowski, D. C. and J. Gruber (2007). Moral hazard in nursing home use. *Journal of Health Economics* 26(3), 560–577.
- Hoynes, H. W. and D. Whitmore Schanzenbach (2009). Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Applied Economics* 1(4), 109–139.
- Kaiser Commission on Medicaid and the Uninsured (2011). Medicaid home and community-based service programs: Data update. Technical report, Henry J. Kaiser Family Foundation.
- Kaiser Commission on Medicaid and the Uninsured (2013, September). Five key facts about the delivery and financing of long-term services and supports.
- Kaiser Commission on Medicaid and the Uninsured (2016). Medicaid home and community-based services programs: 2013 data update. Technical report, Henry J. Kaiser Family Foundation.
- Kaplow, L. (2011). *The theory of taxation and public economics*. Princeton University Press.
- Kemper, P. (1988). The evaluation of the national long term care demonstration: Overview of the findings. *Health Services Research* 23(1), 161.
- LeBlanc, A. J., C. Tonner, and C. Harrington (2001). State Medicaid programs offering personal care services. *Health Care Financing Review* 22(4), 155–174.
- Lepidus Carlson, B., L. Foster, S. B. Dale, and R. Brown (2007). Effects of cash and counseling on personal care and well-being. *Health Services Research* 42(1), 467–487.
- Mankiw, N. G. and M. Weinzierl (2010, February). The optimal taxation of height: A case study of utilitarian income redistribution. *American Economic Journal: Economic Policy* 2(1), 155–76.
- Moffitt, R. (1989). Estimating the value of an in-kind transfer: The case of food stamps. *Econometrica: Journal of the Econometric Society*, 385–409.
- Munro, A. (1992). Self-selection and optimal in-kind transfers. *The Economic Journal* 102(414), 1184–1196.
- Nichols, A. L. and R. J. Zeckhauser (1982). Targeting transfers through restrictions on recipients. *The American Economic Review* 72(2), 372–377.
- Schneider, A., K. Fennel, and P. Keenan (1999, May). Medicaid eligibility for the elderly.

- Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear IV regression. Technical report, NBER Working Paper.
- U.S. Department of Health and Human Services (1992, April). Estimating eligibility for publicly-financed home care: Not a simple task ...
- Whitmore, D. (2002). What are food stamps worth? *Princeton University Working Paper #468*.

Tables and Figures




Figure 1: Equivalent variations and excess burdens of a subsidy

[Equivalent variations and excess burdens of a price subsidy that reduces the after-subsidy price from p_0 to p_1 for individuals with different levels of demand for the subsidized good. The equivalent variation of the subsidy is increasing in the level of demand for the good (individual B 's equivalent variation, the area bounded by the vertices $ABGF$, exceeds individual A 's equivalent variation, the area bounded by the vertices $ABDC$). The excess burdens of the subsidy are independent of the level of demand and instead depend only on the slope. The excess burden of subsidizing individual A 's purchases of the good is the area bounded by the vertices CDE , and the excess burden of subsidizing individual B 's purchases of the good is the area bounded by the vertices FGH .]




Figure 2: CDFs of Formal Care by Treatment Status
 [Data from the Cash and Counseling follow-up survey. Hours of formal home care per week.]




Figure 3: CDF of Formal Care, NLTC

[Data from the 1999 National Long-Term Care Survey. Hours of formal home care per week. Approximately 86 percent of people consume no formal care. Two individuals consumed more than 168 hours of care per week and have been omitted from the figure.]

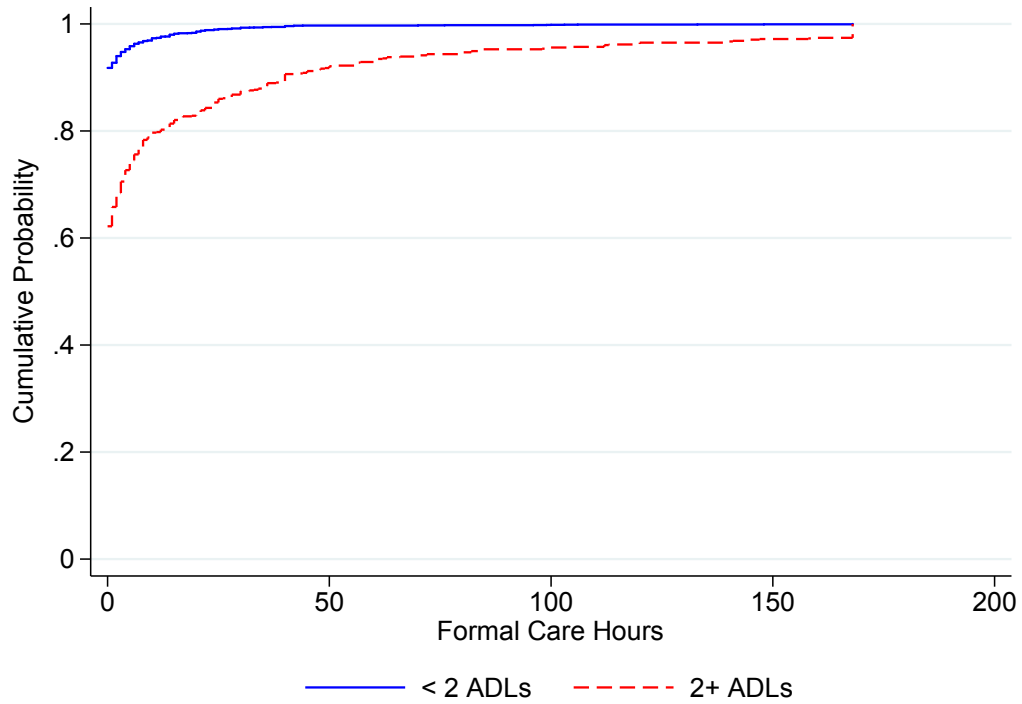


Figure 4: CDF of Formal Care by ADLs

[CDFs for people 65 and older with fewer than vs. at least two activities of daily living limitations. Data from the 1999 National Long-Term Care Survey. Hours of formal care per week. Two individuals reported consuming more than 168 hours of care per week and have been omitted from the figure.]




Figure 5: Distribution of the demand for formal care
[Simulated distribution of formal care satiation points, α , in hours per week. The population is people age 65 and older with at least two activities of daily living limitations. The mean is 21 hours per week.]




Figure 6: Equivalent variation of mixed cash/in-kind program as function of subsidy rate, σ [Equivalent variation of mixed cash/in-kind program as function of the subsidy rate, σ . Programs with larger subsidy rates have smaller cash benefits in order to hold fixed total program spending.]

Table 1: Summary Statistics and Balance Tests

	Arkansas			Florida			New Jersey		
	Cash	In-kind	Difference p-value	Cash	In-kind	Difference p-value	Cash	In-kind	Difference p-value
Formal care hours, baseline	8.61	7.85	0.31	13.16	13.26	0.95	16.10	15.39	0.52
Number unpaid caregivers, baseline	2.12	2.13	0.90	1.93	2.05	0.41	2.02	2.15	0.36
Age	78.73	79.33	0.25	79.00	79.87	0.22	77.63	77.81	0.77
Male	0.15	0.17	0.56	0.19	0.22	0.46	0.19	0.22	0.31
White	0.61	0.63	0.60	0.67	0.72	0.21	0.50	0.57	0.11
Less than high school degree	0.67	0.67	0.96	0.37	0.38	0.78	0.66	0.63	0.44
High school degree	0.29	0.26	0.36	0.44	0.46	0.67	0.19	0.21	0.61
College degree or more	0.03	0.05	0.13	0.17	0.14	0.33	0.11	0.13	0.36
Health	3.25	3.28	0.70	3.12	3.22	0.17	3.13	3.24	0.07
Lives alone	0.33	0.33	0.85	0.26	0.32	0.10	0.32	0.37	0.23
Observations	428	432	.	245	237	.	312	292	.

Means by state and type of transfer. P-value is for test of equality of means across the cash and in-kind groups within the state. Formal care hours and number of unpaid caregivers are for the baseline survey at the time of randomization. Remaining variables are measured at the nine-month follow-up.

Table 2: Summary Statistics for NLTCS

	Mean	Standard Deviation
Formal care hours	3.34	17.53
Number of ADLs	0.76	1.50
Health	2.37	0.94
Age	78.92	7.83
Male	0.35	0.48
Married	0.44	0.50
Number in household	1.84	0.93
Any children	0.80	0.40
Income eligible for Medicaid	0.59	0.49

Means and standard deviations of variables from the 1999 NLTCS Community survey. Number of observations is 5,147 for most variables. Self-rated health is relative to other persons the same age and varies from 1 (excellent) to 4 (poor). Income eligible for Medicaid indicates whether the person's income would qualify her for Medicaid HCBS in her state.

Table 3: Average Hours of Formal Care by Treatment Group

	Cash	In-kind	Difference p-value
Overall	7.11	14.76	0.00
Arkansas	6.94	11.00	0.00
Florida	7.79	19.35	0.00
New Jersey	6.81	16.60	0.00

Means for formal home care hours per week. Cash indicates group received cash transfer; in-kind indicates group received traditional Medicaid home care. P-value for test of equality across groups shown in last column. Rows denote different samples.

Table 4: Sensitivity of Demand for Formal Care, First Stage Estimates

	(1)	(2)
Assigned to cash	8.14***	8.07***
	(0.25)	(0.25)
Controls	No	Yes
F-Statistic	1,066	1,046
Mean market price	13.73	13.73
Adjusted R-squared	0.35	0.37
Observations	1,946	1,946

Dependent variable is the marginal price of formal care. Data are from the Cash and Counseling experiments. Controls described in text are included in column (2). The mean price paid for formal care by the cash group is \$11.02 per hour; for the in-kind group, it is \$2.87 per hour. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table 5: The Sensitivity of the Demand for Formal Care to the Composition of Benefits

	(1)	(2)
Price	-1.85***	-1.82***
	(0.17)	(0.17)
Controls	No	Yes
Mean hours	10.89	10.89
Observations	1,946	1,946

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Columns (1) and (2) are IV Tobits where formal care hours are censored at zero. Controls described in text are included in column (2). Robust standard errors reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Association of Observable Characteristics and Formal Care Hours

	(1) Use Care	(2) OLS Positive Care	(3) Adjusted Hours	(4) Use Care	(5) OLS Positive Care	(6) Adjusted Hours
Age	0.01*** (0.00)	0.37 (0.25)	0.45*** (0.13)	0.01*** (0.00)	0.44 (0.28)	0.45*** (0.13)
Three ADLs	-0.03 (0.05)	17.87** (7.13)	5.34** (2.67)	-0.03 (0.05)	21.28** (8.69)	6.49** (3.00)
Four ADLs	0.09* (0.05)	14.91** (7.13)	8.37** (3.25)	0.08* (0.05)	18.29** (7.85)	8.23** (3.30)
Five ADLs	0.08* (0.04)	31.30*** (7.31)	13.60*** (3.50)	0.06 (0.04)	28.66*** (7.79)	12.16*** (3.61)
Six ADLs	0.13*** (0.05)	39.76*** (8.36)	22.60*** (4.40)	0.13*** (0.05)	43.30*** (9.50)	22.90*** (4.56)
Good health	0.07 (0.07)	-4.06 (12.85)	-1.30 (5.11)	0.09 (0.06)	-4.28 (13.28)	0.21 (5.33)
Fair health	0.15** (0.07)	-10.71 (12.04)	-0.55 (4.96)	0.14** (0.06)	-11.80 (12.00)	1.25 (5.18)
Poor health	0.17** (0.07)	-1.40 (12.63)	3.25 (5.24)	0.18*** (0.06)	-2.40 (13.24)	5.24 (5.71)
Male	-0.06 (0.04)	-12.33** (5.79)	-3.62 (2.46)	-0.04 (0.04)	-10.93* (6.47)	-3.44 (2.65)
Lives alone	0.16*** (0.04)	20.44*** (5.92)	16.53*** (3.44)	0.15*** (0.04)	18.89*** (6.44)	15.29*** (3.55)
Has children	-0.08* (0.05)	7.22 (6.22)	1.80 (3.47)	-0.10** (0.04)	7.09 (7.24)	0.11 (3.56)
Income	-0.09 (0.07)	9.70 (11.57)	-10.62 (6.74)	-0.10 (0.07)	4.14 (11.80)	-10.01 (6.72)
Income squared	0.04* (0.02)	1.48 (2.80)	4.46** (2.19)	0.03* (0.02)	2.47 (2.43)	4.40** (2.16)
State F.E.	No	No	No	Yes	Yes	Yes
Observations	872	331	872	872	331	872

Column (1) is probit where dependent variable is whether person had positive hours of formal care in past week. Column (2) regresses hours of care in past week on covariates for subsample with positive hours. Column (3) uses the entire sample, but adjusts the hours of care used to reflect the price of care in the state. Columns (4) - (6) same as (1) - (3) except that state fixed effects have been included. Income is measured in thousands of 1999 dollars. Omitted categories are 2 ADLs and excellent health. Data are from the 1999 NLTCs. Robust standard errors reported. * p<0.10, ** p<0.05, *** p<0.01

Table 7: Heterogeneity in Demand for Formal Care

	(1) If Positive Care	(2) OLS Positive Care	(3) Adjusted Hours
State fixed effects	0.07	0.11	0.06
Health controls	0.12	0.21	0.13
Informal care options	0.12	0.24	0.15
Income	0.14	0.28	0.17

Each cell reports the R-squared (columns (2) and (3)) or McFadden's Pseudo R-squared (column (1)) of a regression of care hours on the specified set of variables. Column (1) is estimated with a probit; column (2) with OLS conditional on positive care hours; and column (3) with OLS. Each row adds the specified set of variables to the regression. Health controls include age, indicators for number of activities of daily living limitations, indicators for self-rated health, and gender. Informal care options include indicators for the whether person lives alone and whether she has children. Income includes a linear and a quadratic control for the individual's income. Sample restricted to those in 1999 NLTCs with at least two activities of daily living limitations.

Table 8: Summary Statistics by Take-up Decision, NLTC

	(1) Take-up = 0	(2) Take-up = 1	(3) Difference p-value
Age	81.29	82.34	0.27
Four or more ADLs	0.48	0.61	0.01
If health fair or poor	0.64	0.76	0.02
Male	0.32	0.22	0.04
Lives alone	0.25	0.45	0.00
Monthly income	919.42	665.37	0.00

Means presented separately for those who had not taken up Medicaid home care (column (1)) and those who had (column (2)) as well as p-value from test of equality of means across groups. Data from the 1999 NLTC. Only those who had at least two activities of daily living limitations are included. This leads to a sample of 849 individuals, 104 of whom took up. Income is reported in dollars per month.

	(1)	(2)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	
	Baseline	0	β	25	50	0	Max	1	2	3	4	$\gamma = 1$	$\bar{c} = \$2.5k$	Drop $\alpha > 50$	α distribution
Optimal policy															
Subsidy rate, σ^*	0.88	1.00	0.77	-0.50	-0.50	0.94	0.86	1.30	-0.50	0.69	1.10	-0.50	0.89	0.59	0.75
Equivalent variation over pure-cash policy, \$1,000s															
Optimal subsidy policy	5.53	7.60	0.71	0.00	6.09	4.28	>38.13	0.67	1.86	35.74	0.23	21.07	1.74	2.31	
First-best policy	9.38	7.60	1.96	0.22	8.35	9.52	-	-	-	-	-0.08	25.39	2.63	3.34	
Non-care consumption, \$1,000s															
Mean, optimal subsidy	15.83	15.00	20.75	21.87	15.14	15.98	12.89	20.38	16.97	14.30	20.38	15.77	17.30	16.34	
Mean, pure-cash policy	19.56	16.16	21.87	21.87	17.99	21.14	19.56	19.56	19.56	19.56	19.56	19.56	18.61	17.56	
Std. dev., optimal subsidy	1.24	0.00	1.17	0.00	0.65	1.34	3.20	6.03	3.00	1.05	6.03	1.13	1.56	1.22	
Std. dev., pure-cash policy	5.61	5.98	0.13	0.00	5.34	5.84	5.61	5.61	5.61	5.61	5.61	5.61	2.56	3.09	
Consumption distortion															
Total CV over total cost	0.48	0.65	0.88	-	0.44	0.51	0.02	-	0.63	0.26	-	0.52	0.85	0.73	
$E(q_{FC} optimal subsidy)$	14.01	15.80	2.84	0.00	11.67	17.73	20.31	3.78	11.55	17.30	3.78	14.15	5.37	4.97	
$E(q_{FC} pure-cash policy)$	5.72	13.38	0.02	0.00	5.51	6.26	5.72	5.72	5.72	5.72	5.72	5.72	2.33	2.18	
Targeting benefit															
Corr(marg. utility, CV)	0.84	0.90	0.24	-0.68	0.90	0.81	0.11	0.19	0.64	0.85	-0.94	0.81	0.74	0.81	
$E(1 subsidy \succ \text{cash pol.})$	0.16	0.24	0.04	0.00	0.18	0.14	0.14	0.83	0.15	0.15	0.83	0.16	0.18	0.13	

Table 9: Policy analysis and robustness. Column 1 presents results based on the baseline assumptions. Columns 2–5 vary the value of β away from the baseline value of 1.8. Columns 6 and 7 vary the values of the α 's corresponding to people who consume no formal care when facing a positive price (which are only partially-identified). Column 8 sets these α 's to zero. Column 9 sets these α 's to the maximum value consistent with these individuals' choices to consume no formal care when facing a positive price. Columns 10–11 use different models of state-dependent utility in which $\mu(\alpha)$ is linear in α and in which the multiplier factors $\mu(\alpha)$ vary by a factor of 100, $\max_{\alpha}\{\mu(\alpha)\}/\min_{\alpha}\{\mu(\alpha)\} = 100$. In columns 8 and 9, the utility function is "inner state-dependent." In columns 10 and 11, the utility function is "outer state-dependent." In columns 12 and 13, $\mu(\alpha)$ is decreasing, and in columns 14 and 15, $\mu(\alpha)$ is increasing. See Appendix C for more details about state-dependent utility. Column 12 sets the coefficient of relative risk aversion to one (log utility), whereas the baseline coefficient of relative risk aversion is three. Column 13 sets the consumption floor to \$2,500, whereas the baseline value is \$5,000. Column 14 drops values of α (formal care satiation levels) that exceed 50 hours per week. Column 15 cuts every α value in half. 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 units of formal care). "Total CV over cost" is the total ex-post compensating variation of benefits under the optimal program as a fraction of the total cost of these benefits. Mean values of formal care consumption, $E(q_{FC})$, are in hours per week. "Corr(marg. utility, CV)" is the correlation between marginal utility in the absence of any policy and the ex-post compensating variation of benefits under the optimal subsidy. " $E(1|subsidy \succ \text{cash pol.})$ " is the fraction of people who prefer the optimal subsidy to the pure-cash policy benefit ex post.

	Tag: Lives alone		Tag: Number of ADL limitations			
	No	Yes	2-4	5	6	
Average formal care consumption, h/w	9.3	25.1	8.8	19.3	27.3	
Optimal policy, \$s in \$1,000s						
Tagged pure-cash benefits, $(B = b)$	5.67	10.46	6.08	6.4	7.92	
Tagged mixed benefits, (B, σ, b)	(5.67, 0.87, 1.4)	(10.46, 0.9, 1.77)	(6.02, 0.87, 1.52)	(8.53, 0.9, 1.11)	(7.92, 0.92, 0.68)	
Equivalent variation over untagged policy, \$1s						
Tagged pure-cash benefits, $(B = b)$	227		4			
Tagged mixed benefits, (B, σ, b)	44		8			
Targeting benefit						
Corr(marg. utility, tagged pure-cash benefit)	0.20		0.05			

Table 10: Tags analysis. Average formal care consumption, in hours per week, is estimated in the NLTCs. The sample consists of people age 65 and older with at least two activities of daily living limitations. 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 units of formal care). “Corr(marg. utility, tagged pure-cash benefit)” is the correlation between marginal utility in the absence of any policy and the optimal tagged pure-cash benefits.

Appendices

A Theory Appendix

A.1 The optimal mix of in-kind and cash benefits

Consider a planner choosing how to allocate a given budget, B , between cash and in-kind benefits. The planner's goal is to choose the benefits package that maximizes expected utility:

$$\max_{\sigma} EU(\sigma) = \int_{\Theta} v(p(\sigma), m(\sigma, B); \theta) f(\theta) d\theta.$$

The first-order condition, which holds with equality at an interior optimum, σ^* ,³⁴ is

$$\begin{aligned} \frac{dEU(\sigma^*)}{d\sigma} &= \int_{\Theta} \frac{dv(p(\sigma^*), m(\sigma^*, B); \theta)}{d\sigma} f(\theta) d\theta = E_{\Theta} \left(\lambda(\sigma^*; \theta) \frac{dV(\sigma^*; \theta)}{d\sigma} \right) = 0 \\ \iff Cov_{\Theta} [\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0] &= (\sigma^* p_X^0) E_{\Theta} (\lambda(\sigma^*; \theta)) E_{\Theta} \left(\frac{dx_X(\sigma^*; \theta)}{d\sigma} \right). \quad (3) \end{aligned}$$

The second version of Equation 3 shows that, at the margin at an optimum, the covariance between marginal utility and the level of demand for X must be the same sign as the mean marginal change in X due to the shift in benefit composition, i.e., $\text{sign}(Cov_{\Theta} [\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0]) = \text{sign} \left(E_{\Theta} \left(\frac{dx_X(\sigma^*; \theta)}{d\sigma} \right) \right)$. This is the classic equity-efficiency tradeoff. Absent distortions, $E_{\Theta} \left(\frac{dx_X(\sigma^*; \theta)}{d\sigma} \right) = 0$, the optimal benefit composition fully eliminates the covariance between marginal utility and the demand for X , $Cov_{\Theta} [\lambda(\sigma^*; \theta) x_X(\sigma^*; \theta) p_X^0] = 0$. More generally, the greater is the marginal distortion cost of shifting toward in-kind provision, the greater must be the marginal targeting benefit.

The first version of Equation 3 implies that, at the margin at an interior optimum, the benefit to some types from shifting toward greater in-kind provision must be exactly offset by the cost to other types of this shift. Suppose there are just two types, L and H . Then at an interior optimum, the end of the first row of Equation 3 implies that at the margin the planner optimally imposes

$$\left| \frac{dV(\sigma^*; \theta_L)}{d\sigma} \right| = \frac{p_H \lambda_H}{(1 - p_H) \lambda_L} \frac{dV(\sigma^*; \theta_H)}{d\sigma}$$

dollars' worth of costs on L types in exchange for \$1 worth of benefits to H types. The marginal willingness to pay in terms of costs imposed on L s in order to help H s by \$1 is increasing in the ratio of the expected marginal utility of H s to the expected marginal utility

³⁴In certain contexts, including possibly home care, it might be feasible to subsidize formal care at more than a 100 percent rate, so that consumers face a negative net-of-subsidy price of formal care. In this case, the subsidy rate σ can take any real value and the first-order condition holds with equality. A necessary condition for a greater-than-100-percent subsidy to be feasible is that recipients are not able to freely dispose of the good.

of L_s .

A.2 First best

In the first-best case, an individual's type, θ , is verifiable. In this case the planner can choose different (b, σ) benefit bundles for different types. The total derivative of type θ 's indirect utility with respect to the in-kind component of its benefit, σ , is

$$\begin{aligned} \frac{dv(p(\sigma), m(\sigma, B); \theta)}{d\sigma} &= \lambda(\sigma; \theta) \left[x_X(\sigma; \theta)p_X^0 - x_X(\sigma; \theta)p_X^0 - (\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma} \right] \\ &= -\lambda(\sigma; \theta)(\sigma p_X^0) \frac{dx_X(\sigma; \theta)}{d\sigma}, \end{aligned}$$

which is negative for all positive subsidy rates. When type is verifiable, a pure cash contract is optimal, and the cash benefits for each type are chosen to equalize each type's marginal utility. Verifiable types means that the planner can redistribute across types without resorting to distortions, so there is no motive for introducing a distortion in this case.

B Robustness and Generalizability of the Estimate of the Demand for Formal Home Care

As we discuss in Section 6, the key conclusion about the desirability of subsidizing formal care is robust to a wide range of values of the price sensitivity of demand for formal care. But the magnitudes of the optimal subsidy and the welfare gains from in-kind provision depend on the particular value of the price sensitivity of demand. The price sensitivity of demand for care is important for other questions as well, including the extent to which private long-term care insurance contracts that subsidize formal care suffer from a "moral hazard tax." In this section, we address issues related to both the internal and external validity of our estimates of the price sensitivity of demand for formal care.

B.1 Internal validity

There are two main threats to the internal validity of our estimate of the price sensitivity of demand for formal care. The first is quantity constraints that might limit consumption of traditional Medicaid home care. If quantity constraints bind, the first stage of our IV overstates the change in prices (marginal values) associated with being randomized to the cash group and thereby leads us to underestimate the price sensitivity of demand. Quantity constraints may have taken two main forms in this context: supply constraints and statutory or de facto limits on Medicaid home care benefits.

Supply constraints are thought to have faced Medicaid home care recipients in Arkansas during the period of the Cash and Counseling experiment (Brown et al., 2007). These

constraints apparently arose from some combination of Medicaid paying below-market prices and the local home care market being in disequilibrium around the time of the experiment. To the extent that such issues were important, ignoring them would tend to lead us to underestimate the true price sensitivity of demand. The simplest way to avoid this issue is to drop Arkansas from the analysis and instead focus on Florida and New Jersey.

Quantity constraints may also have arisen from statutory or de facto limits on how much Medicaid home care people can use. Both Arkansas and New Jersey had statutory limits on Medicaid home care—16 hours per week in Arkansas and 25 hours per week in New Jersey. (Florida had no statutory limit.) Moreover, as discussed in the text, the amount of Medicaid home care that someone can consume is determined by a care plan written by the individual’s physician. If physicians, whether in an effort to be “good agents” of Medicaid or for other reasons, prescribe care plans whose hours fall short of their patients’ satiation points, then Medicaid home care recipients may not be able to reach their satiation points.

Although in principle the combination of maximum benefit limits and care plan limits could limit the quantity of Medicaid home care available to recipients, in practice it does not appear that either one of these constraints significantly constrained consumption. On care plans, many recipients consume strictly less than their care plan hours, and it is not clear what incentive physicians may have to restrict hours. If anything, physicians’ professional norms and ethos might lead them to act as an agent of the patient rather than Medicaid. Maximum benefit limits also appear to be less binding than might have been expected. LeBlanc et al. (2001) survey Medicaid home care programs and discuss several explicit mechanisms for granting exceptions to the limits. For example, in New Jersey, where the statutory limit was 25 hours per week, with prior authorization a recipient could receive between 26 and 40 hours of care per week and with central office approval a recipient could receive as much care as “needed.” Consistent with these or other mechanisms relaxing quantity limits, the distributions of formal care hours among Cash and Counseling participants receiving traditional Medicaid home care do not exhibit much bunching around these limits. If the limits were binding, one would expect significant bunching because a binding limit causes a convex kink in the budget constraint between formal care and all other goods.³⁵ Figures D.1–D.3 present the CDFs of formal care hours for people randomized to the in-kind group in each of the three Cash and Counseling states. In Arkansas (Figure D.1), there is no apparent bunching that would suggest that consumption was constrained by the state’s limit. In addition to there not being a large mass point at 16 hours, nearly one-fifth of the sample consumed more care than the state’s limit. In New Jersey (Figure D.3), there is bunching at certain points in the CDF of care hours, but this appears to be more of a function of rounding than any limits being imposed. The mass points at 15 and 20 hours (8 and 9 percent of the distribution, respectively) are similarly sized to the mass point at the statutory limit of 25 hours (11 percent).

In Table D.1, we present estimates of the price sensitivity of formal care for each state. The first row shows that the IV Tobit estimates range from -0.96 (Arkansas) to -2.79 (Florida). In the second row, we impose the upper bounds on care hours implied by the Arkansas

³⁵Of course, any test of bunching faces the limitation that measurement error lessens observed bunching. A useful feature of our context in this regard is that the tested-for kink in the budget constraint is quite sharp, from zero up to the market price. To the extent that care limits were truly binding, one might expect the limits to be highly salient to recipients and as a result perhaps less attenuation from reporting error.

and New Jersey limits. We censor observations above those cutoffs and use the IV Tobit to re-estimate the price sensitivity. The additional censoring reduces our estimated price sensitivity in Arkansas but increases it in New Jersey. (We exclude Florida since care hours are not limited there.) The differences across states are similar to those found with the standard IV Tobit.

Generally, the results are consistent with the concern that quantity constraints—whether from supply constraints in Arkansas or statutory limits in Arkansas and New Jersey—might be biasing our price sensitivity estimates towards zero. The state without limits (Florida) consistently displays greater price sensitivity than the other states. Because average care consumption is so different across states, it is also useful to consider the percentage changes implied by the coefficients. A one-dollar increase in the price of formal care is estimated to increase formal care consumption by 9 percent in Arkansas, 14 percent in Florida, and 10 percent in New Jersey. The results also reveal important heterogeneity in price sensitivity across states above and beyond that which appears to be due to quantity constraints. We return to this issue in our discussion of external validity below.

The second main threat to the internal validity of our estimate of the price sensitivity of demand for formal care is the distributional assumptions we make in the estimation. The key assumption we make is that the unobservables are jointly normally distributed (particularly that ε_i , the residual in the latent demand function, is normal). This assumption is important because the majority of the cash group and a large minority of the in-kind group do not consume any formal care. People who do not consume any formal care are at a corner, so revealed preference analysis only bounds their level of demand. The Tobit normality assumption is one way among many to deal with this missing data problem.

We test the sensitivity of our results to a number of different distributional assumptions on ε_i . In each case, we continue to instrument for price as we did in the main analysis. These results can be found in Table D.2. As seen in columns (2) through (4), the estimated price sensitivity changes somewhat from one specification to the next but not dramatically so.

In the next four columns of Table D.2, we assume that everyone who is potentially at a corner solution has a marginal value of care of exactly p , the maximum consistent with their behavior. As seen in Figure 2, those in the cash group were more likely to consume zero hours of care than those in the in-kind group. In a Tobit model, this greater mass at the censoring point tends to reduce the (latent) mean of the care hours distribution for the cash group relative to the in-kind group. The 2SLS model does not have this feature and, as a result, tends to produce smaller mean differences between the cash and in-kind groups. In our setting, this translates into a smaller price sensitivity. Again, we instrument for the price of care with each participant’s randomly assigned transfer type. Under these assumptions, we tend to find a price sensitivity around -1.

As we show in Section 6, only values of the price sensitivity far greater than any we find in this appendix section can overturn the result that the optimal subsidy on formal care is significantly greater than zero.

B.2 External validity

The generalizability of the results from the Cash and Counseling experiments to other contexts depends on the similarity of the experiments' participants to various populations of interest (in terms of price sensitivity of demand for formal care) and how well the experiments match various policies of interest.

Cash and Counseling participants are unlikely to be representative of Americans 65 and older in bad health. Most participants selected into Medicaid home care, and Medicaid home care recipients have a greater demand for formal care than the population as a whole. The participants are also unlikely to be representative of the broader population of Medicaid home care recipients. Participation in the Cash and Counseling demonstrations is voluntary and the benefits are increasing in the price sensitivity of demand for formal care. By participating, an individual gains the possibility of receiving in cash roughly the cost to Medicaid of providing their formal care benefit. The extent to which an individual values the cash benefit more than the in-kind benefit is increasing in the sensitivity of the individual's demand for formal care to its price. It is natural to expect that participants in the experiments were more sensitive to the price of formal care than the broader population of Medicaid home care recipients in the Cash and Counseling states. This tends to increase our estimate of the price sensitivity of demand for formal care relative to what we would expect to find among the broader population of recipients of Medicaid home care.

Another reason the results of the Cash and Counseling experiments might not generalize well to other contexts is the nature of the experiment itself. Care-giving arrangements, for which people often make important investments such as moving or adjusting their labor supply, likely depend on both the past history of policies and expectations about future policies. People arrange their lives in order to make the best of the choices available to them, and their decisions about where to live and work and whether to use formal or informal home care likely depend on the nature of any home care benefits for which they might be eligible. The Cash and Counseling experiments likely came as a surprise to many participants, and it is unclear what participants might have expected about the persistence of this policy—would it continue indefinitely or would they soon be reverted back to traditional Medicaid home care? Both the surprise aspect and the uncertainty about how long cash benefits might last likely dampened responses relative to what they would have been under an anticipated, permanent policy.

These considerations suggest caution in applying the results of the Cash and Counseling experiments to other contexts. But the robustness of our main conclusions to even large changes in the price sensitivity of demand for formal care greatly limit this concern in our context. And the strengths of the Cash and Counseling experiments—the large, exogenous price variation—make it a valuable piece of evidence about the demand for formal care and the effects of alternative home care-related policies.

C Policy Analysis

C.1 Estimating the distribution of demand for formal care

As discussed in the text, we use the observed distribution of formal care consumption together with our estimate of the price sensitivity of demand for formal care to infer the latent distribution of the level of demand for formal care. We express the level of demand for formal care in terms of satiation points, α . The only tricky part of this calculation is that observed formal care consumption does not point-identify α for people consuming zero formal care, it only bounds it: $\alpha_i \leq \beta p_i$. We estimate the full α distribution, including the α 's of people who consume zero formal care, in three steps.

The first step involves using the observed distribution of formal care consumption, q , to infer the partially-unobserved distribution of latent demand, q^* , where $q_i = \max\{0, q_i^*\}$. In the baseline specification, we fill in the censored values of q_i^* corresponding to the $q_i = 0$ cases by linearly extrapolating the observed q density among people with small positive quantities. In particular, we calculate the number of people in each of two groups: those who consume more than zero and less than five hours of care per week and those who consume more than five and less than ten hours of care per week. Based on the shares of people in each group, we estimate the implied (constant) slope of the probability density function over this range and its level at $q^* = 0$. We assume that this slope remains constant at lower values of q^* , which amounts to assuming that the left part of the underlying latent quantity distribution has a triangular distribution. For each censored q^* (corresponding to an individual who consumed no formal care at market prices), we draw the underlying latent q^* from the truncated triangle distribution based on the estimated slope. Figure D.4 shows the underlying distribution of formal care consumption on which this calculation is based.

Second, we convert each q^* to its corresponding α using the estimated price sensitivity of demand for formal care, $\alpha_i = q_i^*(p) + \beta p$. This adjusts (potentially latent) formal care consumption by our estimate of the impact of the price on consumption. Finally, we estimate the kernel density of the implied α distribution. Figure 5 shows the resulting α distribution. It is mostly just a rightward-shifted version of the observed distribution of formal care consumption, with adjustments for the censoring of people who consume no formal care.

For the tags analysis, we repeat the same procedure for estimating the α distribution separately for different groups of people, as defined by their tagged characteristics. Figures D.5 and D.6 show the α distributions of people who do vs. do not live alone and for people with different numbers of activities of daily living limitations. All of the distributions are similarly-shaped, and they exhibit the expected differences in levels. The demand for formal care is greater among people who live alone than among people who live with others, and it is greater among people with more activities of daily living limitations.

We test the robustness of our results to making different extreme assumptions about how to fill in the unidentified α values. In one case, we set every unidentified α value to zero, which is equivalent to assuming that anyone who consumed no care when facing market prices would also consume no care when facing a price of zero. In the other extreme, we set all of the partially-identified α 's equal to their (point-identified) upper bound, $\alpha_i = \hat{\beta} p_i$.

C.2 State-dependent utility

As discussed in the text, any state-dependence in utility that is correlated with formal care consumption is centrally important for the value of in-kind provision, since it affects the value of redistribution across people with different levels of demand for formal care. State-dependence that increases the marginal utility of people with greater demand for formal care relative to people with lower demand for formal care increases the attractiveness of in-kind formal care transfers, whereas state-dependence that decreases the marginal utility of people with greater demand for formal care relative to people with lower demand for formal care decreases the attractiveness of in-kind formal care transfers. Given the possibility that people with different demands for formal care might have systematically different utility functions, it is therefore important to test the robustness of the results to different possibilities about state-dependent utility.³⁶

Two natural ways in which to model state-dependent utility are to introduce a scaling factor on the outside or inside of the utility function:

$$U(c; \theta) = \begin{cases} \mu(\theta)u(c), & \text{“outer state-dependence”;} \\ u(\mu(\theta)c), & \text{“inner state-dependence”}. \end{cases}$$

“Outer state-dependence” multiplies the standard, type-independent component of the utility function by a factor $\mu(\theta) \geq 0$, which is potentially correlated with demand for formal care. This type of state dependence has a straightforward effect on the value of redistribution across types. Types with greater scaling factors have greater marginal utility for any given level of net consumption. “Inner state-dependence” multiplies net consumption (non-care consumption net of any utility costs of residual health problems) inside the standard, type-independent utility function. Unlike “outer state-dependence,” “inner state-dependence” can have a subtle effect on the marginal utility of a given level of net consumption. On the one hand, types with greater scaling factors are more effective at converting income into net consumption (“effective consumption” is $\mu(\theta)c$, which is increasing in $\mu(\theta)$ for any c), which tends to increase the marginal utility of income. On the other hand, types with greater scaling factors have greater effective consumption for any given level of net consumption, which tends to reduce the marginal utility of income due to marginal utility diminishing in the level of effective net consumption. With log utility, these two effects cancel out, and “inner state-dependence” has no effect on the marginal utility of income. With preferences in which marginal utility diminishes more rapidly in effective consumption, such as constant relative risk aversion preferences with a coefficient of risk aversion greater than one, the latter effect dominates and types with greater scaling factors have lower marginal utility for any given level of net consumption.

³⁶Although health-dependent utility is a natural concern, in the context of home care benefits its importance is somewhat diminished by the fact that most home care benefit programs limit eligibility to people with at least two activities of daily living limitations. This ensures that home care benefits go only to people who have fairly severe chronic health problems. As a result, the type of state-dependence of utility that is relevant for the design of home care benefits (taking as given the eligibility criteria for home care benefits) is state-dependence within the set of (sick) people eligible for benefits, not between people in good vs. bad health.

D Appendix Figures and Tables




Figure D.1: CDF of Formal Care in Cash and Counseling States, Arkansas
[Data from the Cash and Counseling follow-up survey of the in-kind group in Arkansas. Formal care is measured in hours per week. Arkansas had a regulation that limited care to 16 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]




Figure D.2: CDF of Formal Care in Cash and Counseling States, Florida
 [Data from the Cash and Counseling follow-up survey of the in-kind group in Florida. Formal care is measured in hours per week. Florida had no regulation limiting care hours (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]




Figure D.3: CDF of Formal Care in Cash and Counseling States, New Jersey
 [Data from the Cash and Counseling follow-up survey of the in-kind group in New Jersey. Formal care is measured in hours per week. New Jersey had a regulation that limited care to 25 hours per week (LeBlanc et al., 2001). The vertical dotted lines mark 10, 15, 20, and 25 hours per week for reference.]




Figure D.4: Distribution of formal care consumption among people with two or more ADL limitations

[Distribution of formal care consumption among people with two or more activity of daily living limitations in the NLTCs. The figure omits the 65 percent of people who report consuming no formal care and the 3 percent of people who report consuming more than 150 hours per week of formal care for readability. The mean of the full distribution is 12 hours per week.]




Figure D.5: Distribution of demand for formal care by whether someone lives alone
 [Estimated probability density functions of formal care satiation points, α , for each of two groups: people who do not live alone (left-most pdf) and people who do live alone (right-most pdf). The mean of the distribution is 16 hours per week among people who do not live alone and 37 hours per week among people who do live alone.]




Figure D.6: Distribution of demand for formal care by number of ADL limitations
 [Estimated probability density functions of formal care satiation points, α , for each of three groups: people with 2-4 ADL limitations (left-most pdf), people with five ADL limitations (middle pdf), and people with six ADL limitations (right-most pdf). The mean of the distribution is 16 hours per week among people with 2-4 ADL limitations, 31 hours per week among people with 5 ADL limitations, and 34 hours per week among people with six ADL limitations.]

Table D.1: Sensitivity of Demand and Statutory Limits

	(1) Arkansas	(2) Florida	(3) New Jersey
Price, IV Tobit	-0.96*** (0.25)	-2.79*** (0.46)	-1.71*** (0.15)
Price, IV Tobit Limits	-0.45*** (0.12)		-1.93*** (0.16)
Controls	Yes	Yes	Yes
Market price, formal care	12.36	15.09	14.59
Mean hours, in-kind group	11.00	19.35	16.60
Observations	860	482	604

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Separate regressions run for each state with IV Tobit (first row). Second row uses IV Tobit and imposes statutory limits as upper bounds on care hours. Controls described in text are included in all regressions. Robust standard errors reported. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.2: The Sensitivity of the Demand for Formal Care to the Composition of Benefits

	Censored errors			Uncensored errors			
	Normal	Extreme Value	Logistic	T-location scale	Normal	Negative binomial	Poisson
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Price	-1.85*** (0.14)	-2.47*** (0.24)	-1.40*** (0.10)	-1.21*** (0.08)	-0.94*** (0.09)	-0.72*** (0.11)	-1.07*** (0.21)
Mean hours	10.89	10.89	10.89	10.89	10.89	10.89	10.89
Observations	1,946	1,946	1,946	1,946	1,946	1,946	1,946

Dependent variable is hours of formal care per week. Data are from the Cash and Counseling experiments. Columns (1) - (4) are IV specifications where the error term is treated as censored on the left. Each column presents the estimated sensitivity of demand under a different distributional assumption on the underlying error term. Columns (5)-(7) use distributions that implicitly assume there is no censoring on the left. All models instrument for price with the participant's randomized treatment status and are estimated via two-stage residual inclusion. for columns (6) and (7), average marginal effects are reported. * p<0.10, ** p<0.05, *** p<0.01