

Costs and Benefits of In-Kind Transfers: The Case of Medicaid Home Care Benefits

Ethan M.J. Lieber and Lee M. Lockwood



Costs and Benefits of In-Kind Transfers: The Case of Medicaid Home Care Benefits

Ethan M.J. Lieber

University of Chicago

Lee M. Lockwood

Northwestern University

October 2013

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

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 5 RRC08098401-05-00). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Regents of the University of Michigan

Mark J. Bernstein, Ann Arbor; Julia Donovan Darlow, Ann Arbor; Laurence B. Deitch, Bloomfield Hills; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park ; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Costs and Benefits of In-Kind Transfers: The Case of Medicaid Home Care Benefits

Abstract

Many large government programs provide benefits in kind as opposed to in cash. Providing benefits in kind potentially distorts decisions and leads to a deadweight loss if recipients value the benefits less than a cost-equivalent cash transfer. Yet providing benefits in kind may have some offsetting benefits, especially in terms of improving the targeting of benefits to desired beneficiaries. We complete what is to our knowledge one of the first empirical studies of the costs and benefits of providing transfers in kind as opposed to in cash. We focus on the case of the US Medicaid program's provision of in kind home health care benefits. Three state Medicaid programs completed randomized experiments that converted the usual in-kind benefits into cash benefits for a randomly-selected subset of benefit recipients. We use the results of these experiments together with a variety of other evidence to estimate the costs and benefits of providing Medicaid home care benefits in kind. We find that in the case of Medicaid home care benefits, both the costs and benefits of providing transfers in kind as opposed to in cash are large. This suggests that alternative targeting mechanisms, if available, have the potential to significantly increase efficiency relative to traditional Medicaid policy.

Citation

Lieber, Ethan M.J. and Lee M. Lockwood (2013). "Costs and Benefits of In-Kind Transfers: The Case of Medicaid Home Care Benefits." Ann Arbor MI: University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2013-294.
<http://www.mrrc.isr.umich.edu/publications/papers/pdf/wp294.pdf>

Authors' Acknowledgements

We are grateful for helpful comments from Seema Jayachandran and participants at the Michigan Retirement Research Center Workshop. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) through the Michigan Retirement Research Center (Grant #5 RRC08098401-05), funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of the Social Security Administration, any agency of the Federal Government, or the Michigan Retirement Research Center.

1 Introduction

Economists have long considered the question of whether transfers should be provided in cash or in-kind. A fundamental principle of economic theory suggests that consumers will, in general, at least weakly prefer receiving transfers in cash to cost-equivalent in-kind transfers. Cash transfers afford recipients greater flexibility in tailoring their consumption to their tastes and circumstances. This is a major advantage of cash over in-kind transfers. Yet providing transfers in-kind potentially has offsetting benefits, such as improving the targeting of benefits to intended recipients, reducing moral hazard in the context of the Samaritan's Dilemma, and increasing the efficiency of the tax system. Thus, it is an empirical question whether a particular transfer is more efficiently provided in kind or in cash.

The question of the mode of the transfer is not purely academic. In-kind transfers are ubiquitous in government programs, charities, and private insurance. Currie and Gahvari (2008) report that early in the 21st century, the United States government spent more than 12 percent of GDP on just three in-kind programs combined (health, child care, and education).

Much of the empirical literature to date has focused on the costs of in-kind transfer programs (the consumption distortion), and much of this literature has focused on food stamp programs (e.g., Moffitt, 1989; Whitmore Schanzenbach, 2002). Much of the theoretical literature, by contrast, has considered possible benefits of in-kind transfers, such as improved targeting of benefits to desired beneficiaries (Nichols and Zeckhauser, 1982; Blackorby and Donaldson, 1988), improved tax system efficiency (Munro, 1992), and the Samaritan's dilemma (Bruce and Waldman, 1991). Such benefits imply that providing benefits in kind can in some cases be optimal, despite the resulting consumption distortion.

Prominent among the potential benefits of providing in-kind benefits is improved targeting efficiency. Governments and private insurance pools face a major difficulty: Information problems limit their ability to finely condition the benefits they deliver on the state of the world, i.e., to effectively target benefits to intended beneficiaries and not others. In practice they have to make transfers contingent on coarse bundles of states rather than particular states, and the coarse bundles may include component states with widely varying levels of optimal transfers. Well-chosen in-kind benefits can potentially alleviate this problem. Because there is heterogeneity in the values that people place on a particular good, a planner could strategically choose the bundle of goods that is provided in-kind to attract only a subset of the population to participate in the program (targeting on the extensive margin) or lead to differential intensity of program participation among participants

(targeting on the intensive margin). Providing transfers in cash does not automatically allow for the same targeting benefits: a dollar is worth a dollar to everyone. Providing transfers in cash makes the program equally attractive to both targets and non-targets.

In this paper, we undertake a theory-guided investigation of the costs and benefits of providing transfers in-kind. Our empirical application is in health care, a large and rapidly-growing part of government budgets in many countries. In particular, we focus on the provision of home health care benefits by the US Medicaid program. Medicaid provides certain individuals assistance with tasks to enable them to live at home instead of a nursing home. For example, some beneficiaries can obtain help with bathing, housekeeping, eating, and other personal care tasks. Medicaid contracts with approved agencies to provide this “formal care” to recipients. However, Medicaid participants supplement their “formal care” with “informal care” provided by family and friends. Home health care is a particularly good setting to assess the costs and benefits of in-kind provision for a number of reasons. First, informal care is used widely and can be a close substitute for formal care in many situations. That implies the consumption distortion due to in-kind provision could be very large. On the other hand, there is considerable unobservable heterogeneity in the costs of informal care across individuals. That implies that in-kind provision could lead to significant targeting benefits. To quantify the trade-offs facing both social and private insurance, we (i) measure the consumption distortion from the in-kind provision of Medicaid home care benefits using randomized experiments; (ii) provide a variety of evidence on the targeting benefits of the in-kind provision of Medicaid home care benefits; and (iii) discuss the implications for Medicaid and other programs.

To estimate the consumption distortion due to in-kind provision of transfers, we use data from experiments that the Arkansas, Florida, and New Jersey Medicaid programs implemented. In these experiments, participants were randomized to receive traditional Medicaid benefits (control group) or the value of the benefits in cash (treatment group). Those who were assigned to the treatment group received the cash transfer, but were only allowed to use the money for home health care. Although the cash was restricted in its use, it allowed Medicaid recipients to not only use regular formal care, but to pay their informal care providers. Thus, being randomized to treatment greatly expanded the set of home health care goods the recipient could purchase. We find that on average, members of the treatment group reduced their use of formal care by more than 50 percent. This suggests that Medicaid recipients’ home health care use is severely distorted by the in-kind provision of the transfer. We then estimate the deadweight loss due to this consumption distortion.¹ We find that the deadweight loss due to in-kind provision ranges from \$526 to

¹As discussed in much more detail in the paper, the experiments do not directly estimate the demand curve necessary to calculate the deadweight loss due to the consumption distortion. To adjust the demand

\$2,505 per Medicaid home care user per year. Extrapolating those figures to the entire Medicaid home health care population, the aggregate deadweight loss due to the consumption distortion ranges from \$1.5 billion to \$7.1 billion per year.

Given our finding of a significant consumption distortion cost, it is important to evaluate the potential benefits of providing home care in-kind. Currie and Gahvari (2008) review the large body of theory on targeting and conclude that providing transfers in-kind does not dramatically increase targeting efficiency because there is often a means test or other barrier to participation in the program. Despite this assessment, we estimate the extent to which providing the current bundle of benefits in-kind helps Medicaid target its “desired beneficiaries,” defined in various ways. We find that paying benefits in kind offers a significant advantage in terms of targeting efficiency compared to paying cash benefits. In other words, Medicaid’s particular bundle of in-kind transfers helps target benefits toward desired beneficiaries. Our estimates suggest that the targeting benefits due to the in-kind provision of Medicaid home care are approximately \$1.7 billion.

We do not find strong evidence that providing benefits in-kind increased tax system efficiency or reduced moral hazard in the context of the Samaritan’s Dilemma. Taken together, our results suggest that the benefits of in-kind provision are near the bottom end of our estimated range for the cost of the consumption distortion.

2 Medicaid personal care services and the Cash and Counseling experiments

Medicaid provides home and community-based services (HCBS) primarily through two programs: the Medicaid Title XIX PCS optional State plan and the Medicaid 1915(c) HCBS waiver program.² Typical care services include help with housekeeping, bathing, dressing, and eating. Medicaid HCBS have grown rapidly in recent years. In 1999, 1.9 million people received care through the HCBS programs; by 2007, that number had expanded to 2.8 million. In addition to the growing number of participants, the fraction of Medicaid long-term care dollars that go to HCBS has risen from 19 percent in 1995 to 42 percent in 2008 (Kaiser Commission on Medicaid and the Uninsured, 2011). Because of rising use and the economic downturn Medicaid HCBS budgets have been under

curve that comes out of the experiments, we use variation in minimum wage laws that shifts the costs of home health care in the general population. This instrumental variables strategy recovers an estimate of the elasticity of demand for formal care that, when combined with the experimental results, allows us to estimate the deadweight loss due to the consumption distortion.

²More detailed summaries of Medicaid-provided personal care services are available in LeBlanc et al. (2001) and Kaiser Commission on Medicaid and the Uninsured (2011).

tremendous pressure to keep costs down. This has led to a number of demonstration projects that tested different approaches to reducing the costs of providing these services. We use data from experiments run in Arkansas, Florida, and New Jersey. They are collectively known as the Cash and Counseling experiments (C&C).

Although the three states differed in their exact implementations, each experiment had a few basic features. After completing a baseline survey, consumers were randomized to either receive services as usual (control group) or to direct their own personal care (treatment group). Treatment group members were given a monthly stipend with which they could hire workers or purchase other services or goods for their personal care needs (e.g., assistive devices, home modifications). However, treated members had to keep receipts for their purchases to show that they were spending the money on personal care services.³ Thus, treatment group members did not simply receive a straight cash transfer. Instead, they received a cash benefit that could be spent on a specific class of goods.

Personal care services can be classified as either formal or informal care. Formal care is provided by home health or personal care agencies while informal care is usually provided by family and friends. Medicaid programs traditionally provide formal care (in-kind benefit). To receive formal care, a Medicaid recipient creates a care plan with a physician or registered nurse and then has that plan approved by a designated agency (e.g., the Arkansas Foundation for Medical Care in Arkansas). Conditional on approval, the physician or registered nurse and the consumer choose a provider from a list of Medicaid-qualified agencies. In the experiments, treatment group members were not restricted to hire care providers from the list of approved agencies. This greatly expanded their choice set for personal care.

In Arkansas and New Jersey, the cash allowance given to treatment group members was not simply the full value of the care in their care plans. In these states, Medicaid recipients had previously used only a fraction of the care outlined in their care plans.⁴ Thus cashing out the full value of the care plan would have compensated the treatment group members by more than the value of the care they could have expected to receive under the regular Medicaid program. In Florida, analysis showed that Medicaid recipients were receiving, on average, all the services described in their care plans. Consequently, there was not an adjustment in the number of hours that were cashed out for treatment group members. However, all three states cashed out care plans at less than their full value to pay the

³Treated consumers could spend up to 10 percent of their allowance on services that could not be readily invoiced (e.g., a neighbor mows the lawn).

⁴They do not receive the full extent of the plan for a number of reasons, the most frequent being institutionalization (e.g., a nursing home) and the home care workers simply not showing up to deliver the specified care.

counselors who helped treatment group members construct their plans.⁵ To the extent that treatment group members were not fully compensated for the value of care they could receive under traditional Medicaid (all treatment group members had the option of leaving the experiment at any time and resuming their normal Medicaid care), substitution between formal and informal care will be understated.

Although consumers were randomized to the treatment and control groups (50 percent chance of being treated), not every participant completed the nine-month followup survey. Averaged across the three states, 20 and 35 percent of the treatment and control groups, respectively, attrited. If attrition is correlated with being in the treatment group and with care use, then the exogeneity of treatment could be compromised. However, as seen in the top panel of Table 1, there does not appear to be any differential attrition across the two groups. Baseline use of care, health, gender, age, race, and education are all very similar across the sample of treated and control group consumers. Of the 30 balance tests shown in the top panel, none of the differences between treatment and control groups are statistically distinguishable from zero at the 5 percent level.

The bottom panel of Table 1 reports the hours of formal and informal care the consumers received in the two weeks preceding the nine-month followup survey. This comparison reveals that on average, treatment group members consumed about 8 fewer hours of formal care per week in Arkansas, 23 fewer hours in New Jersey, and 18 fewer hours in Florida. These differences suggest that in-kind provision of home care may cause a substantial consumption distortion.

3 The cost of in-kind benefits: consumption distortion

The main cost of providing benefits in kind is a deadweight loss from distorting consumption; receiving benefits in kind may lead some people to consume more of that good than they otherwise would have had they received a cash benefit. Other things equal, recipients would be better off receiving a cost-equivalent cash transfer instead. Figure 1 shows this graphically. The figure depicts an in-kind benefit program that provides up to two units of formal care to eligible recipients. This shifts the budget line two units outward in the direction of greater formal care. Unlike a cost-equivalent cash transfer, however, the in-kind transfer does not increase the maximum amount of other goods (other than formal care) that recipients can consume at low levels of formal care. For recipients who would

⁵The experimental demonstrations had budget neutrality requirements that necessitated such a step.

have chosen to consume at least two units of formal care had they received a cost-equivalent cash transfer, the in-kind transfer is equivalent to a cash transfer and there is no consumption distortion cost. For recipients who would consume less than two units of formal care if they received a cost-equivalent cash transfer, the in-kind transfer is less valuable than a cost-equivalent cash transfer. The difference between the cost of the transfer and the smaller value they place on it is a deadweight loss from providing the transfer in kind as opposed to in cash. In the figure, this loss is the distance between the hypothetical cash-transfer budget line (dashed black line running from (0,7) to (7,0)) and a parallel budget line that would allow the individual to achieve the utility level she achieves under the in-kind transfer program. In the case of the individual with blue indifference curves, this cost is slightly less than the cost of one unit of all other goods (or, equivalently, slightly less than the cost of one unit of formal care).

When trying to estimate the magnitude of the consumption distortion caused by the in-kind provision of transfers, researchers have generally focused on food stamp, or direct food transfer, programs (Cunha, 2013; Hidrobo et al., 2012; Hoynes and Whitmore Schanzenbach, 2009; Moffitt, 1989; Whitmore Schanzenbach, 2002). Moffitt (1989) first estimated the marginal value of food stamps using a 1982 policy change in Puerto Rico which converted food stamps into direct cash payments. He found that there was very little distortion induced by in-kind provision. The majority of consumers were inframarginal, i.e., the transfers were small enough that most people purchased food with cash (in addition to that purchased with the stamps) and so faced the full price for their marginal purchases. He then argues that for the minority of consumers whose consumption was distorted, a trafficking market made the food stamps worth very close to their face value in cash. Whitmore Schanzenbach (2002) finds similar results to Moffitt (1989) using experiments in San Diego's and Alabama's food stamp programs. The experiments randomized participants to receive food stamps or the cash value of the stamps. Although the majority of consumers were inframarginal, those who were not valued food stamps at approximately eighty cents on the dollar.

Instead of focusing directly on the marginal value of the good provided in-kind, we are primarily interested in estimating the welfare loss due to the consumption distortion.⁶ To measure this welfare loss, we need to estimate the demand curve for formal care. With that demand curve, we can then calculate the deadweight loss triangle associated with any overconsumption of formal care. The Cash and Counseling Experiments effectively generate a large, discrete jump in the price of formal care. Under traditional Medicaid, the

⁶There is also a related literature on whether cash transfers should be restricted or unrestricted (Baird et al., 2011). The C&C program only varies whether the transfer was provided in-kind, not whether there were restrictions on the transfer.

dollar price of formal home care is zero – Medicaid provides the care in-kind. Although the dollar price is zero, the shadow price of Medicaid care that incorporates the difficulty of scheduling with providers and other such costs is certainly positive. We focus our discussion on the dollar price for care, but recognize it is just an approximation to the full price the consumer faces. For consumers in the treatment group, the marginal dollar price for formal home care jumps up to the market price. Based on the Gemworth Financial surveys of home care agencies, the average price for an hour of home care in Arkansas, New Jersey, and Florida was \$12.72 (\$10.90 in Arkansas, \$12.65 in Florida, and \$14.60 in New Jersey).⁷

Table 2 reports the average hours of formal care along with the estimated percentage impact of treatment. The first column pools all three states and finds that treatment reduced the number of formal care hours by just over 50 percent. The next three columns present the same results separately for each state. In each state, treatment group members used considerably less formal care than control group members. The estimates suggest that on average, raising the dollar price of formal care by \$1 will reduce the use of formal care by 3.4 percent in Arkansas, 4.1 percent in Florida, and 3.4 percent in New Jersey. On average, a \$1 increase in the price of formal care reduces the quantity used by 3.6 percent.

Although the experiment’s randomization provides an exogenous source of variation, it does not directly recover the demand curve necessary to estimate the deadweight loss due to the consumption distortion. Instead, it provides an upper bound on this deadweight loss. Figure 2 shows a Medicaid home care user’s demand curve for formal care. When Medicaid provides care in-kind, consumers pay nothing out-of-pocket for it and so value that care somewhere between \$0 and the market price at that quantity. Thus, consumers are initially at a point like *A*. However, those treated in the C&C experiment face the market price for care and locate at the equilibrium point *B*. The experimental evidence essentially connects points *A* and *B* to estimate the demand curve. Unfortunately, this overstates the demand curve’s steepness which affects the deadweight loss calculation. The true deadweight loss is given by *C*, the area between the true demand and supply curves past the market equilibrium. An estimate of deadweight loss based on the C&C demand curve counts not only *C*, but *D* as well.

To solve this problem, we estimate the elasticity of demand for home care in the population more generally and use that estimate to recover the price that C&C participants would have been willing to pay for the quantity of formal care they received under Medicaid. As seen in Figure 2, with this price p_1 , we can estimate the actual

⁷The prices reported in the survey were from 2005. Because the experiments happened around 2000, we adjust the prices not only by the CPI, but also for price growth in the home care sector. For Arkansas, we observe the price that Medicaid paid for care in 2000 and use that price, combined with the Gemworth Financial survey data to make the adjustment.

deadweight loss due to the consumption distortion, not just an upper bound on that deadweight loss. Let ε_{pop} be the elasticity of demand for the population. And let $\varepsilon_{cc}(d\ln(q), p_2, p_1)$ be the elasticity for people in the C&C experiments as a function of the percentage change in quantity, and the old (p_1) and new (p_2) price for care. If those two elasticities are equal, $\varepsilon_{pop} = \varepsilon_{cc}(d\ln(q), p_2, p_1)$, and we know ε_{pop} , $d\ln(q)$, and p_2 , then we can recover p_1 .⁸ Note that both p_2 and $d\ln(q)$ are available from our previous analysis of the C&C data (\$12.72 and -53 respectively). The final piece needed is the elasticity of demand for home care. We now turn to estimating it.

We use changes in states' minimum wage laws to estimate the demand curve for formal care in the population more broadly. In particular, changes in minimum wage laws shift the costs of providing formal home care and identify points along the demand curve. The underlying identifying assumption is that minimum wage laws shift the supply curve around without directly affecting the demand curve, i.e. that the only reason minimum wage laws impact the quantity demanded is through their impact on the price of formal care.

First, we test whether the minimum wage laws affected how much home health care workers are paid. Using data from the Bureau of Labor Statistics for the fifty states and Washington D.C., we estimate the following equation.

$$wages_{st} = min\ wage\ amount_{st}\beta + X_{st}\Gamma + \lambda_s + \lambda_y + \varepsilon_{st} \quad (1)$$

where *wages* is the average hourly wage for home health care workers in each state and year, *min wage amount* is the relevant minimum wage in the state and year, X are demographic including variables for age, gender, and race, λ_s are state fixed effects, λ_y are year fixed effects, and ε is the error term.⁹ Regressions are weighted by state population and standard errors are clustered at the state level.

Table 3 reports the results from equation (1). In the first column, demographic characteristics are omitted. The result indicates that increasing the minimum wage by one dollar will increase home care workers' wages by approximately 41 cents. The coefficient is highly significant and implies a first-stage F-statistic of more than 30. The next column includes the demographic variables mentioned above. The coefficient decreases slightly in magnitude, though it remains highly statistically significant. In the final column of the table, leads of the minimum wage have been included in the regression. Neither is large in magnitude nor anywhere approaching statistical significance. This suggests that the

⁸More specifically, with values for ε_{pop} , $d\ln(q)$, and p_2 , the following equation can be solved to recover p_1 : $(p_2 - p_1)/((p_2 + p_1)/2) - d\ln(q)/\varepsilon_{pop} = 0$.

⁹The age, gender, and race variables indicate the fraction of the population of a given age, gender, and race. We use two age categories, below or above 60, two gender categories, and three race categories, white, black, and other.

minimum wage laws are affecting the wage levels and not simply picking up underlying trends.

The second stage of this analysis is to test whether minimum wage laws affect the number of people employed as home health care aides. There is a long literature that estimates the impact of minimum wage laws on employment more generally. Recently, Meer and West (2012) have studied the dynamic effects of minimum wages on employment and found that an increase in the minimum wage does not lead to job destruction, but a slowdown in the growth rate of new jobs. This suggests that changes in the minimum wage will not have immediate effects, but will affect employment over time. Because of this finding, we estimate the impact of lagged minimum wage laws on the current stock of home care employees. Our specification is

$$\frac{\text{home care employees}_{st}}{1k \text{ population } 60+_{st}} = \text{min wage amount}_{s,t-1}\alpha + X_{st}\Gamma_2 + \lambda_s + \lambda_y + \nu_{st} \quad (2)$$

where the dependent variable is the number of home care employees per 1,000 people 60 or older and the right hand side variables are the same as those described in equation (1) with the exception of the lagged minimum wage. Ideally, we would have a measure of hours of formal care consumed, but, to the extent of our knowledge, those data do not exist.

Instead, the number of home care employees proxies for the equilibrium quantity of care consumed.¹⁰ The counts of home care employees come from the County Business Patterns data.¹¹ Because the number of employees is a stock, there is likely autocorrelation in the error terms, ν_{st} . As such, we assume an AR(1) process and estimate equation (2).

Table 4 presents the results. The baseline specification without demographics implies that a \$1 increase in the minimum wage last year reduces the number of home health care aides per 1,000 people 60 or older by approximately .8. Relative to the mean, 12.8, this implies a reduction of just over 6 percent. When demographic controls are included, this estimate rises to 1.08 in magnitude.

To test whether the changes in minimum wage laws are are simply picking up underlying trends, we include two leads of the minimum wage in equation (2). These results are reported in the final column of Table 4. None of the leads are systematically related to the dependent variable at conventional significance levels. This suggests that we are not simply picking up trends in employment and lends credence to the identifying assumption. Using the identifying assumption, we can divide the impacts reported in Table 4 by those reported in Table 3 to recover the elasticity of demand. Our estimates imply an elasticity of -2.1.

¹⁰Our dependent variable is expressed as a rate so that its units are similar to those in the C&C analysis.

¹¹We used SIC code 8080 and NAICS code 621610 in the appropriate years.

With an estimate for the elasticity of demand in the broader population, we now have all of the pieces needed to solve for p_1 , consumers' willingness to pay for the quantity of formal care received from Medicaid. We find that consumers would have been willing to pay \$9.89 per unit of formal care at the quantity of care provided by Medicaid. The following deadweight loss estimates are subject to a number of caveats: we use different populations to estimate the elasticities that we assume are equal, formal care provided by Medicaid might not be exactly the same good as formal care purchased in the general market, and we use evidence from relatively small changes in the minimum wage and use them to extrapolate to large changes. Subject to these caveats, we estimate the deadweight loss triangle, C in Figure 2, with the formula, $(p_2 - p_1)\Delta q/2$. Our estimates imply that per person-year, the deadweight loss is approximately \$526. To the extent that the supply of formal care is not perfectly elastically supplied and to the extent that the demand for formal care is more elastic in the general population than the Medicaid population, our deadweight loss calculation represents a lower bound on the true value.¹² We can also estimate an upper bound on the deadweight loss, by assuming that participants valued the marginal unit of Medicaid-provided formal care at \$0. We estimate an upper bound on the loss of approximately \$2,505 per Medicaid home care user per year.

The Kaiser Commission on Medicaid and the Uninsured (2011) estimates that 2.82 million people received home care services from Medicaid in 2007. Extrapolating our deadweight loss estimates to this population implies that the consumption distortion due to the in-kind provision of goods is between 1.5 and 7.1 billion dollars per year (in year 2000 dollars).

4 The Benefits of In-Kind Transfers: Targeting Benefits

Providing transfers in kind can improve targeting efficiency by leading to systematic self-selection into program participation by the people who have the most to gain from the program. When the planner wishes to redistribute resources across types that she cannot verifiably distinguish from each other, she can potentially use in-kind benefits to induce a favorable self-selection into program participation if different types place different values on the good(s) in the in-kind benefits bundle.

¹²First, if the supply curve is not in fact perfectly elastic, then the change in price used in the DWL calculation understates the base of the triangle. And second, the demand elasticity estimated from the broader population might be greater than the elasticity of demand for Medicaid participants because Medicaid home care users are likely to be sicker than the general population. If true, our estimated willingness to pay of \$9.89 will be greater than the actual willingness to pay.

Planner’s problem.— Suppose the population is comprised of individuals who potentially differ in terms of their required care, η , their informal care costs, σ , and their labor productivity, α . In the absence of government intervention, consumption is increasing in labor productivity and is decreasing in required care and in care costs. If private arrangements fail to break the link between individual characteristics and consumption,¹³ a utilitarian (or other type of) planner may wish to intervene to redistribute resources from individuals with higher labor productivity and lower care costs to individuals with lower labor productivity and higher care costs.¹⁴

First best: full insurance.— If the planner can verifiably distinguish between different types, then she can achieve the first-best, which involves full insurance. Government transfers exactly offset the disadvantages of having low labor productivity or high care requirements or costs and every type consumes the same amount.

Second best: incomplete information.— If on the other hand the planner cannot verifiably distinguish between different types, then she will in most cases be unable to provide full insurance. How completely the planner is able to insure “type risk” and which tools will prove useful to her in doing so depends on the particular information structure she faces.

Suppose, in the interest of approximating the apparent constraints on many social insurance programs such as the Medicaid program in the US, that the planner is either unwilling or unable to condition transfers on labor productivity, informal care costs, informal care use, and consumption. Instead, the planner is willing and able to condition transfers only on care needs, savings, and formal care use. In this case, the key constraint facing the planner is the inability to condition transfers on labor productivity and informal care costs. This constraint limits the extent to which the planner will be able to provide insurance. This constraint also, as noted by Nichols and Zeckhauser (1982) and others, expands the set of “tools” the planner might wish to use in trying to provide insurance, from traditional cash transfers that impose no constraints on recipients to various types of transfers that come with strings attached, such as means tests and in-kind transfers. The cost of imposing restrictions on recipients is that this makes participation in the program less valuable to them. The benefit is that, whereas the planner cannot transfer different amounts of cash to different types that she cannot verifiably distinguish (since everyone

¹³Two reasons private arrangements may fail to fully insure these risks are (i) some of these risks might be realized before people can enter into risk-sharing contracts, e.g., labor productivity may be largely fixed before people enter adulthood, and (ii) some of these characteristics might be unverifiable, i.e., it may not be possible to write contracts that condition on the values of certain characteristics perhaps due to information problems or legal restrictions.

¹⁴Depending on the nature of the process by which people come to have these characteristics, the planner’s intervention could be considered social insurance (if one’s characteristics are realized as the result of a risky process, including perhaps one’s endowment at birth) or redistribution.

values cash and can simply report that they are of the type due the largest cash benefit), the planner can potentially make different transfers to types she cannot distinguish by imposing well-chosen conditions on transfer recipients that are least costly to those types she wishes to target.

Figure 3 shows a simple example of how this can work in practice. Providing benefits in kind, if combined with imposing some other restriction on beneficiaries behavior, can lead certain types to self-select into the program at higher rates than others. In the figure, the additional “restriction” (in addition to providing formal care in kind when resale is impossible) is simply modeled as a fixed “buy in” to participate in the program, paid in one unit of “all other goods”. Paying this cost entitles recipients to two units of formal care, free of other charges. This “buy in” cost can be thought of as a stand in for the utility cost of any restrictions that the program might impose on beneficiaries. State Medicaid programs, for example, limit eligibility to people with low levels of income and assets. These can be thought of as restrictions on the choices that would-be recipients can make if they want to participate in the program.

In the figure, individuals who would purchase at least two units of formal care even in the absence of the program (like those represented by the green indifference curves) are strictly better off participating in the program and thereby having their purchases of formal care effectively subsidized. Individuals who would purchase little if any formal care in the absence of the program (like those represented by the blue indifference curves) may be better off not participating in the program. Although they may place some value on the in-kind formal care benefits provided by the program, it may not be sufficiently large to justify paying the other costs of participating in the program, whether a financial “buy in” or a binding restriction on behavior such as labor supply and earned income. Clearly, pure-cash programs, i.e., programs in which both the costs of participating and the benefits bundle are comprised exclusively of cash, cannot induce such a systematic self-selection into program participation. In-kind programs, by contrast, in which either the costs of participating or the benefits bundle are comprised at least partly of goods, can potentially lead to a systematic self-selection into program participation. The type of goods that have the biggest potential to induce differential selection into program participation of different groups are those for which the demand is much greater among some groups than others and of which resale is costly.

In the following sections we try a variety of approaches to estimate the targeting properties of the bundle of in-kind home-care benefits currently provided by many state Medicaid programs. The goal is to measure the extent to which benefits are received by various groups that a planner might wish to target to greater or lesser extents, and to compare this

allocation to that which might be expected to arise under a means-tested cash benefit program. This comparison should reveal how providing benefits in kind affects the distribution of benefits received by different groups.

The analysis proceeds in two parts. In the first, we measure the targeting effects of providing benefits in-kind on the extensive margin (i.e., whether someone participates in Medicaid) using the Health and Retirement Study. In the second part, we measure the targeting effects of providing benefits in-kind on the intensive margin (i.e., benefits received among Medicaid participants) using data from Cash and Counseling experiment.

For a variety of possible reasons, the planner might place a higher value on a given increment to wealth among some people (“targets”) than among others. Paying benefits in kind as opposed to in cash potentially helps the planner make larger net transfers to the (unverifiable) types she wishes to target. This improvement in targeting efficiency can be achieved if offering the particular bundle of in-kind benefits leads those types that the planner wishes to target to self-select into the program at higher rates than non-targets or to participate more intensively in the program than non-targets. The targeting benefit of paying transfers in kind is positively related to the extent to which paying transfers in kind shifts resources from non-targets to targets.

Measuring the improvement, if any, in targeting efficiency from paying transfers in kind as opposed to in cash faces two main challenges. First, it requires knowledge of the distribution of benefits under both the in-kind and the counterfactual in-cash benefits programs. For simplicity, in our main comparisons we assume that take up of pure cash benefits would be 100%. This is a natural benchmark, as a pure cash program does not involve any cost of claiming benefits, whether in terms of hassles, stigma, or other sources. Moreover, it is straightforward to adjust the results to situations in which take up of the (impure) counterfactual cash benefit program is less than perfect.¹⁵

The second and greater challenge in measuring the targeting benefits of in-kind benefits is translating any given difference in the distribution of net benefits between the in-kind program and the counterfactual cash program into a single-value measure of the change in targeting efficiency. Fully accomplishing this requires specifying a social welfare function and the underlying utility functions of each individual. At the other extreme, one could simply characterize the distributions of net benefits under both types of programs and not attempt to translate the change in the distribution into any single-valued change in social welfare. We take a middle road. The main thrust of our analysis is characterizing the distributions of net benefits under both types of programs. Then, given these distributions,

¹⁵See Currie and Gahvari (2008) for a review of take up of many programs. Ganong and Liebman (2013) document take up of the US food stamps program and find that it has fluctuated significantly over time.

we make rough calculations of the effects on social welfare for a variety of social welfare functions.

The key task of the empirical work is therefore to measure the extent of redistribution toward various groups who differ in the extent to which they are targeted, which depends on the planner’s objective function. In selecting groups across which to measure the extent of redistribution, we focus on characteristics that affect the feasible rate of consumption. In the simplest model in which everyone shares the same preferences, this approach is consistent with the approach of a utilitarian planner, who wishes to transfer resources from people with high rates of consumption (and so low marginal utilities of income) to people with low rates of consumption (and so high marginal utilities of income).¹⁶ In what follows, we measure redistribution across groups of people who differ in their consumption opportunities due to underlying differences in permanent income, health, and informal care costs.

4.1 Targeting on the extensive margin: Evidence from the Health and Retirement Study

The Health and Retirement Study (HRS) is a representative panel of Americans age 51 and over. The HRS has surveyed participants on a variety of information about themselves roughly every two years, with the longest-tracked cohort first interviewed in 1992. It contains especially rich information about respondents’ health and wealth. The key variables for our analysis are: an indicator variable for whether the individual was covered by Medicaid;¹⁷ variables that determine an individual’s eligibility for Medicaid benefits, including income and wealth; and variables that are associated with plausible definitions of “targetedness”, such as proxies for permanent income (e.g., years of schooling, total income during retirement, or Social Security income) and proxies for informal care costs (e.g., whether the individual has a spouse, the health status of the individual’s spouse, and how many children the individual has).

Table 5 presents Medicaid participation rates among several different groups, where

¹⁶The assumption that everyone has the same preferences, although standard in much of the optimal policy literature, is potentially problematic in the context of health problems, since health might affect the mapping between consumption and utility. See, e.g., Viscusi and Evans (1990) and Finkelstein et al. (2013) on health-dependent utility.

¹⁷Respondents are asked whether they are covered by Medicaid. We interpret this to mean that the respondent has received at least some medical care that was paid for at least partially by Medicaid over the period since the last interview. Unfortunately, details about the exact nature of the services covered by Medicaid or their dollar value are not available. Fortunately, such information is available in the C&C data, which we analyze next.

individuals are assigned to groups on the basis of the values of characteristics related to their “targetedness,” i.e., the likely value that a utilitarian planner would place on marginal changes in their wealth. The figure reveals that more targeted individuals—those with lower permanent income, in worse health, or with higher informal care costs—are more likely to participate in Medicaid. This holds both in the full population and in the subpopulation of people whose income and wealth likely makes them eligible for Medicaid.

Similarly, OLS and probit regressions (not reported) of Medicaid participation on various measures of targetedness and demographic variables indicate that, holding other characteristics fixed, increasing an individual’s targetedness on one dimension (e.g., reducing permanent income) is associated with greater Medicaid participation. In both the bivariate and multivariate analyses, the association between Medicaid participation and measures of targetedness tend to be strongest for the proxies for permanent income and health status and are somewhat weaker for the proxies for informal care costs.

Overall, the results suggest that Medicaid redistributes resources toward groups that a utilitarian planner is likely to wish to target—those with low permanent income, bad health, and costly informal care options. Moreover, the resulting redistribution to these targeted groups is greater than that which would be achieved by a cash-benefit program. In other words, providing transfers in kind as opposed to in cash appears to increase the extent of redistribution to targeted types. This occurs both because members of targeted groups are much more likely to be eligible for Medicaid than members of non-targeted groups and because, conditional on being eligible for Medicaid, members of targeted groups are much more likely to participate in Medicaid than members of non-targeted groups.

4.2 Targeting on the intensive margin: Evidence from the Cash and Counseling experiment

We use the C&C data to assess whether in-kind provision provides any targeting benefits on the intensive margin, i.e. for those who do receive Medicaid home care assistance. The data contain measures of the recipients’ health, permanent income/ability, and price of informal care. We begin by comparing whether the targeted group used more Medicaid home care at baseline than the untargeted groups. At baseline, everyone received benefits in-kind. In this instance, correlations between care use and our measures of targetedness are preferable to treatment effects; so long as the unobservable, true measure that the program would like to target on is positively correlated with our measures, then the correlation allows our measure to stand in as a proxy for the underlying set of unobservables that would ideally be used for targeting—a treatment effect for our measure

would lose these correlations with the unobservable components.

Table 6 presents the mean hours of Medicaid provided formal care used at baseline for the targeted and untargeted groups. Only those who were enrolled in Medicaid's home care programs at baseline and those who did not purchase formal care in the private market are included. We excluded the latter group because we cannot separate Medicaid provided formal hours from those purchased in the private market. However, less than 10 percent of the sample is excluded for this reason and including them does not qualitatively change the results.

Our first measure indicates whether or not the person lived alone. People who live alone are likely to face a higher price for informal care (because there is no one living with them who can help) and so would tend to be a targeted group. However, at baseline, those who lived alone actually received fewer hours of formal care from Medicaid than did those who did not live alone (the difference is statistically significant with $p < 0.001$). We might think that people who live alone are actually in better health than those who do not live alone. In that case, if Medicaid targets sicker people, the observed difference in means could be consistent with targeting benefits on the intensive margin. However, as seen in the next rows of Table 6, people who are in better health relative to others their age do not receive systematically lower hours of care than those in poorer health. This suggests that, on the intensive margin, the in-kind provision of care is having little targeting benefit on the health or informal care price dimensions. A third dimension of targeting we can explore with the C&C data is ability. Education levels are frequently used as a proxy for permanent income or ability. As seen in Table 6, there is not a strong relationship between education levels and the use of formal care at baseline. Thus, on the intensive margin, there is no evidence that Medicaid is targeting specific groups or that the provision of care in-kind is allocating greater benefits to targeted individuals.

Now that we have examined both the intensive and extensive margins, we can perform a simple calculation to estimate the value of the targeting benefits. The targeting benefits of in-kind provision occur on the two dimensions discussed. On the extensive margin, in-kind provision reduces participation from 100 percent down to 54 percent of eligibles.¹⁸ On the intensive margin, the fraction of benefits that participants use falls from 100 percent with a cash transfer down to 91 percent with in-kind transfers.¹⁹ The additional cost of Medicaid home care due to the change on the extensive margin is just the full cost of providing home

¹⁸In the HRS, we estimate who is eligible to participate in Medicaid based on income, assets, and ADL limitations.

¹⁹The 90 percent figure is based on data from Arkansas's C&C evaluation. At baseline, recipients had on average been receiving approximately 9.1 hours of formal care per week. Their care plans from Medicaid called for 10 hours of care per week on average.

care multiplied by the fraction increase in participation. This benefit is approximately \$8.9 billion in year 2000 dollars. On the intensive margin, an additional 9 percent of the costs of providing home care to all eligibles is an additional \$2.5 billion in year 2000 dollars. The actual benefit from targeting is not the reduction in the size of the program, but the reduction in tax rates and thereby deadweight loss. If the cost to society of raising a dollar to pay for Medicaid home care is \$0.15 of deadweight loss (Kleven and Kreiner, 2003), then the full targeting benefits are approximately \$1.7 billion in year 2000 dollars.

Of course, our estimate comes with many caveats. First, this calculation assumes that under the in-kind transfers regime, all targets receive the transfer. To the extent this is not true, our estimate overstates the targeting benefits of in-kind provision. Second, our estimate of who is eligible, and therefore of how large the change in participation will be, is inexact. If our measure counts people who are not eligible for Medicaid as eligible, then our participation estimate is too low, which causes our estimate of the increase in participation with a switch to cash transfers to be too large. This would again cause our estimated targeting benefits to be too large.

4.2.1 Testing necessary conditions for targeting benefits in the Cash and Counseling experiment

Nichols and Zeckhauser (1982) discuss a variety of necessary conditions for restrictions imposed on benefits recipients to improve targeting efficiency. Three key conditions are:

- Targeted groups have a greater demand for the good being provided in-kind than do other groups.
- Goods provided in kind must be over-provided relative to what members of targeted groups would have chosen under a budget-equivalent cash transfer.
- Goods provided in kind must not be of higher quality than members of targeted groups would buy under a budget-equivalent cash transfer program.

In this section, we test whether these conditions are satisfied by Medicaid home care benefits.

The first condition is that the targeted group has a greater demand for the good being provided in-kind than does the untargeted group. To test this condition, we compare the use of formal care at followup of targeted and untargeted individuals, but only among participants assigned to treatment. We focus on treated individuals because their use of

formal care is not constrained by Medicaid and therefore provides a clearer picture of their actual demand. Table 7 presents basic summary statistics on the average hours of formal care used in the previous two weeks by treatment group participants. Within each dimension of targeting available in the C&C, there is little evidence that the more targeted group has a higher demand for formal care. Neither living alone nor having poor health are strongly correlated with using more formal care. In fact, those with greater ability appear to have the highest demand for formal care, the opposite of the condition necessary for targeting benefits to have benefits.²⁰

A second necessary condition for in-kind provision of a good to have targeting benefits is that the good must be over-provided relative to what targets would have chosen under a budget-equivalent cash transfer. We test this condition in the C&C data by estimating whether being randomized to treatment reduced targeted participants' hours of formal care. The results are presented in Table 8. Regardless of which dimension of targeting is considered, being assigned to the treatment group is associated with a large reduction in the hours of formal care used. This suggests that Medicaid does in fact over-provide formal care, even for targeted individuals.

A third necessary condition for in-kind provision to have targeting benefits is that Medicaid not provide a higher quality bundle than a target would buy under a budget-equivalent cash transfer. At the followup survey, C&C participants were asked whether they were satisfied with the care they were receiving. Table 9 presents the fraction of participants that were satisfied with their care broken out by treatment status and level of targetedness. For many of the groups we think of as targeted, those who received treatment expressed a (statistically significant) higher level of satisfaction with their care than did those who did not receive treatment. This is consistent with Medicaid providing a relatively low-quality bundle of formal care which is in turn consistent with the necessary condition for in-kind provision to have targeting benefits.

Overall, the in-kind provision of formal care passes two of the three tests that are necessary for targeting benefits to exist. Although the failure of the first test could be interpreted to reject the possibility that in-kind provision could even potentially have targeting benefits in this context, we would urge caution in taking this interpretation because our measures of who is a targeted individual may not correspond exactly to those that Medicaid is actually trying to target.

²⁰One interpretation of these data is that in order for those with more education to qualify for Medicaid, they must have received a much worse draw from the health distribution. However, that is not consistent with the participants' views of their own health.

5 The Benefits of In-Kind Transfers: Samaritan's Dilemma and Tax System Efficiency

Providing home-care benefits in kind has at least two other potential benefits relative to providing benefits in cash. First, it might alleviate the Samaritan's dilemma (Bruce and Waldman, 1991). The Samaritan's dilemma says that the possibility of receiving public welfare or private charity in the future may induce moral hazard in potential recipients thus making it more likely that they do indeed receive welfare or charity. This concern can potentially justify providing certain benefits in kind. Bruce and Waldman (1991), for example, argues that the Samaritan's dilemma suggests that it may be optimal to provide job training rather than cash support since job training may reduce recipients' future reliance on welfare or charity. Similarly, providing home care benefits in kind could alleviate the Samaritan's dilemma. By providing home care benefits in kind, Medicaid likely reduces the prevalence of informal care among the low-income elderly. This informal care is often provided by the adult children of the elderly and often at the expense of engaging in market activity. Supplying informal care often reduces people's market-related human capital (and so their future earnings) in addition to reducing their current earnings. Both of these effects may increase the likelihood that informal caregivers will later rely on Medicaid and other means-tested programs themselves, in their own old age. By providing home care benefits in kind and thereby reducing the demand for informal care, it is possible that Medicaid reduces the demand for future means-tested benefits. Hoynes et al. (2012) suggest that a similar intergenerational mechanism is relevant in the case of the food stamp program.

Although we cannot directly assess the long-term impacts on informal caregivers with the C&C experiments, we can estimate whether switching to cash transfers increased the number of hours of care provided by informal caregivers. Table 10 presents the mean hours of informal care that C&C participants had in the two weeks prior to the followup survey. On average, members of the treatment group had 118 hours of informal care while members of the control group had 112. Even though we cannot reject that these two numbers are statistically distinguishable from each other, suppose that treatment did increase the use of informal care by six hours in a two-week period. The question then is whether three additional hours per week of informal care, a 4.9 percent increase, will significantly impact the probability that the informal caregiver takes up Medicaid home care (or other welfare programs) in the future. Offsetting the potential increase in hours of informal care is the payments that are now received for a portion of those hours (treatment group members were able to pay their informal caregivers with the money from Medicaid). Foster et al. (2007) show that informal caregivers for people in the C&C

treatment group report less financial strain, were less likely to miss work, and were more able to look for other jobs than were informal caregivers for people in the control group (who received formal care in-kind). These findings suggest that even if there is a slight increase in informal caregiving when benefits are paid in-kind, it is unlikely that it causes increased reliance on welfare programs in the future.

Providing home care benefits in-kind could also improve tax system efficiency. A well-known and fundamental principle of public finance is that taxing all goods at equal rates tends to minimize the excess burden from taxation. Yet as noted by (Munro, 1992) and recently emphasized by Mirrlees et al. (2010), when certain goods cannot be taxed or are not taxed (such as non-market time), then tax system efficiency can be enhanced by deviating from the uniform-taxation principle in certain specific ways. In particular, reducing tax rates on complements for market work effort and increasing tax rates on substitutes for market work effort can reduce the excess burden from taxation by reducing the distortion between market and non-market activity from the tax system. Mirrlees et al. (2010) argues that this consideration justifies reducing tax rates on market child care (a complement with market activity, especially among mothers of young children), despite the additional complexity this introduces into the tax system. The case of formal home care is very similar. In both cases, the tax system favors informal, household production of care over market care. And both market child care and market formal home care are complements for market activity. This suggests that there might be benefits of reducing tax rates on formal home care and other efforts to encourage formal care rather than informal care. By paying benefits in terms of in-kind formal home care, Medicaid tends to increase the use of formal care at the expense of informal care and thus may go some way toward increasing tax efficiency.

Although the C&C experiments might appear to reduce tax system efficiency because they encourage additional hours of home care provided by informal caregivers, a fraction of those hours are now taxed as market work. Using the C&C data, Carlson et al. (2007) show that the treatment group members consumed the same or slightly more hours of paid care than control group members. Given that the use of formal care providers fell by more than 50 percent, this means that treatment group members chose to pay previously informal caregivers for the hours they no longer received from formal caregivers. Carlson et al. (2007) also show that the number of informal care hours that were provided without any direct cash payment was actually lower in the treatment group than in the control group. Providing transfers in cash reduced the amount of untaxed informal care provided and either did not affect or actually increased the hours of taxed care provided. Thus, counterintuitively, switching away from in-kind provision of home care to restricted cash payments, if anything, increased tax system efficiency.

6 Discussion of results and conclusion

In this paper, we estimate the costs and benefits of providing Medicaid home health care benefits in kind as opposed to in cash. We find that both the costs (in terms of distorting consumption) and the benefits (in terms of targeting intended beneficiaries) are significant. The finding of significant benefits means that, absent other means of targeting transfers to intended beneficiaries, providing benefits in kind has significant targeting benefits for the Medicaid program. The finding of significant costs, however, indicates that these targeting gains do not come cheaply and that the gains from finding alternative means of targeting benefits to intended beneficiaries would be significant.

There are many difficulties involved in quantifying the costs and benefits of providing transfers in kind, and for this reason our paper should be viewed as a first step in this direction. The issue of optimal benefits 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. The costs and benefits of the design of these programs must be evaluated on a case-by-case basis, and our hope is that our approach in this paper will prove fruitful in the analysis of many of these other policies as well.

Our results have implications for various reform proposals under consideration that seek to increase the control that beneficiaries have over the use of the benefits they receive from various programs. There have been a variety of proposals for Medicaid and other programs to increase the extent of “consumer-driven” care. Many of these proposals effectively bring the in-kind benefits programs closer to cash programs. The primary impetus for these proposals is the view that the consumption distortion from in-kind benefits is large, a view that is supported by our analysis of Medicaid home care benefits. Yet an overlooked cost of such reform proposals is that, in addition to reducing the consumption distortion cost of in-kind benefits, such proposals also likely reduce the targeting benefits of in-kind benefits. These reductions in targeting benefits must therefore be set against the reductions in consumption distortions in determining optimal policy. This is an important topic for future research.

References

- Baird, S., C. McIntosh, and B. Özler (2011). Cash or condition? evidence from a cash transfer experiment. *The Quarterly Journal of Economics* 126(4), 1709–1753.
- Blackorby, C. and D. Donaldson (1988). Cash versus kind, self-selection, and efficient transfers. *The American Economic Review*, 691–700.
- Bruce, N. and M. Waldman (1991). Transfers in kind: Why they can be efficient and nonpaternalistic. *The American Economic Review* 81(5), 1345–1351.
- Carlson, B. L., 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.
- Cunha, J. M. (2013). Testing paternalism: Cash vs. in-kind transfers. *American Economic Journal: Applied Economics* forthcoming.
- 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.
- Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association* 11(s1), 221–258.
- Foster, L., S. B. Dale, and R. Brown (2007). How caregivers and workers fared in cash and counseling. *Health services research* 42(1p2), 510–532.
- Ganong, P. and J. B. Liebman (2013). The decline, rebound, and further rise in snap enrollment: Disentangling business cycle fluctuations and policy changes. Technical report, National Bureau of Economic Research.
- Hidrobo, M., J. Hoddinott, A. Margolies, V. Moreira, and A. Peterman (2012). Impact evaluation of cash, food vouchers, and food transfers among colombian refugees and poor ecuadorians in carchi and sucumbíos. *Washington, DC: International Food Policy Research Institute*.
- Hoynes, H. W., D. W. Schanzenbach, and D. Almond (2012). Long run impacts of childhood access to the safety net. Technical report, National Bureau of Economic Research.
- 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*, 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.
- Kleven, H. and C. Kreiner (2003). The marginal cost of public funds in oecd countries; hours of work versus labor force participation.
- LeBlanc, A. J., C. Tonner, and C. Harrington (2001). State medicaid programs offering personal care services. *Health Care Financing Review* 22(4), 155–174.
- Meer, J. and J. West (2012, June). Effects of the minimum wage on employment dynamics. *SSRN*, 1–28.
- Mirrlees, J. A. et al. (2010). *Dimensions of tax design: the Mirrlees review*. Oxford University Press.
- 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.
- Viscusi, W. K. and W. N. Evans (1990). Utility functions that depend on health status: estimates and economic implications. *The American Economic Review*, 353–374.
- Whitmore Schanzenbach, D. (2002). What are food stamps worth? *Princeton University Working Paper #468*.

Figure 1: Consumption Distortion from In-Kind Provision




Figure 2: Market for Formal Care




Figure 3: Targeting Benefits from In-Kind Provision




Table 1: Summary Statistics and Balance Tests

	Baseline characteristics								
	Arkansas			Florida			New Jersey		
	Treated	Control	p-value for difference	Treated	Control	p-value for difference	Treated	Control	p-value for difference
Formal care hours	8.89	8.79	0.89	12.63	13.02	0.76	15.77	15.84	0.94
Number unpaid caregivers	4.35	4.19	0.38	3.90	4.15	0.51	4.05	4.14	0.69
Health	3.19	3.24	0.25	3.11	3.02	0.18	3.21	3.18	0.50
Lived alone	0.34	0.33	0.72	0.27	0.22	0.21	0.35	0.40	0.15
White	0.62	0.65	0.38	0.65	0.51	0.14	0.48	0.54	0.10
Male	0.17	0.17	0.91	0.19	0.22	0.29	0.18	0.23	0.07
Age	77.84	78.03	0.69	77.11	78.03	0.19	75.56	75.71	0.81
Education:									
Less than high school	0.67	0.66	0.76	0.34	0.38	0.24	0.64	0.63	0.70
High school	0.28	0.27	0.57	0.49	0.47	0.63	0.21	0.22	0.53
College or more	0.03	0.05	0.10	0.15	0.13	0.47	0.10	0.11	0.53
	Care use at followup survey								
	Arkansas			Florida			New Jersey		
	Treated	Control	p-value for difference	Treated	Control	p-value for difference	Treated	Control	p-value for difference
Formal care hours	12.38	21.43	0.00	16.09	39.24	0.00	14.32	32.79	0.00
Informal care hours	114.23	110.20	0.55	125.29	122.54	0.77	117.25	107.44	0.24

Sample means presented separately for treated and control groups. 1,237 observations in Arkansas, 688 in Florida, and 858 in New Jersey. Formal care paid for directly by Medicaid at baseline. Informal care hours at baseline are for the one week preceding the survey. Formal and informal care hours at followup are for the two weeks preceding the survey.

Table 2: How Cash Benefits Affect the Use of Agency Care

	Baseline	Separately by state		
		kansas	Florida	New Jersey
Treatment group	13.90 (0.65)	12.38 (0.96)	16.09 (1.40)	14.32 (1.15)
Control group	29.29 (1.15)	21.43 (1.39)	39.24 (3.41)	32.79 (1.39)
Implied percent effect	-52.54** (2.91)	-42.23** (5.85)	-59.01** (5.03)	-56.34** (3.97)
Price change due to treatment	12.71	10.90	12.65	14.60
Observations	2,783	1,237	688	858

The mean of formal care hours by treatment status are presented in the first two rows (standard errors in parentheses). The implied percentage impact of being in the treatment group reported in third row. Baseline column includes observations from all three states that conducted experiments. Price change due to treatment gives the estimated price increase for formal care that corresponds to being treated.

Observations gives the total (treatment plus control) number of observations for the column. * $p < .05$, ** $p < .01$

B

Table 3: Impact of Minimum Wage Laws on Health Aides' Wages

	Add		
	Baseline	demographics	Add leads
Minimum wage	0.406** (0.073)	0.357** (0.084)	0.354** (0.071)
Leads of minimum wage:			
One year			0.009 (0.044)
Two years			-0.051 (0.048)
State f.e.	x	x	x
Year f.e.	x	x	x
Demographics		x	x
Mean of dependent variable	8.928	8.928	8.928
F-stat	30.51	18.07	24.93
R-squared	0.941	0.948	0.948
N	656	656	656

Dependent variable is average wages for home health care aides. Data for 1999-2011. Minimum wage is the minimum wage in effect in a given state and year. Demographics include fraction of population of given gender, age, and race. Standard errors clustered by state. * $p < .05$, ** $p < .01$.

Table 4: Impact of Minimum Wage Laws on Health Aide Employment

	Baseline	Add demographics	Add leads
Minimum wage	-0.838 (0.481)	-1.076* (0.499)	-1.118* (0.507)
Leads of minimum wage:			
One year			-0.058 (0.373)
Two years			-0.466 (0.351)
State f.e.	x	x	x
Year f.e.	x	x	x
Demographics		x	x
Mean of dependent variable	12.819	12.819	12.819
R-squared	0.489	0.532	0.536
N	1020	1020	1020

Dependent variable is home health care aides per 1000 60+ year olds.

Data from 1989-2007. Minimum wage is the minimum wage in effect in a given state in year t-1. Demographics include fraction of population of given gender, age, and race. * p<.05, ** p<.01.

Table 5: Targeting Benefits in HRS, Extensive Margin

	Fraction in Medicaid	Observations
Lived alone	0.45	888
Did not live alone	0.32	1,493
Health:		
Poor	0.48	449
Fair	0.42	835
Good	0.3	656
Very Good	0.25	350
Excellent	0.24	87
Education:		
Less than high school	0.45	1,212
GED	0.31	147
High School	0.29	611
Some college	0.26	299
College or more	0.23	109

Fraction of eligibles who have Medicaid and number of observations in HRS data. Within each dimension of targeting, rows progress from more targeted to less targeted.

Table 6: Targeting Benefits in C&C, Intensive Margin

	Average hours of formal care	Observations
Lived alone	12.19	781
Did not live alone	14.06	1,534
Health:		
Poor	13.18	1,023
Fair	13.74	822
Good	12.98	355
Excellent	14.42	100
Education:		
Less than high school	13.71	1,301
High school	12.32	739
College or more	15.35	211

Means and number of observations presented for the number of formal hours in the previous week at baseline. Within each dimension of targeting, rows progress from more targeted to less targeted. Statistics presented only for those enrolled in Medicaid home care at baseline and were not purchasing non-family care on the private market.

Table 7: Demand for Formal Care by Dimensions of Targeting

	Average hours of formal care	Observations
Lived alone	13.95	460
Did not live alone	13.89	959
Health:		
Poor	13.57	623
Fair	14.31	481
Good	13.78	237
Excellent	12.22	67
Education:		
Less than high school	13.30	823
High school	14.77	439
College or more	16.73	117

Means and number of observations presented for the number of formal hours among treatment group at followup. Within each dimension of targeting, rows progress from more targeted to less targeted.

Table 8: Over-provision of Medicaid Formal Care

	Treated	Control	p-value of difference
Lived alone	13.95	21.70	0.000
Did not live alone	13.89	33.29	0.000
Health:			
Poor	13.57	29.19	0.000
Fair	14.31	29.27	0.000
Good	13.78	29.75	0.000
Excellent	12.22	31.35	0.012
Education:			
Less than high school	13.30	26.24	0.000
High school	14.77	32.65	0.000
College or more	16.73	36.37	0.000

Means of hours of formal care in two weeks preceding followup survey presented by targeting dimension and treatment status. P-value for difference within targeting dimension, across treatment status presented in last column.

Table 9: Quality of Medicaid Formal Care

	Treated	Control	p-value of difference
Lived alone	0.89	0.79	0.000
Did not live alone	0.92	0.89	0.018
Health:			
Poor	0.90	0.82	0.000
Fair	0.91	0.86	0.037
Good	0.94	0.90	0.139
Excellent	0.91	0.91	0.907
Education:			
Less than high school	0.91	0.86	0.002
High school	0.92	0.85	0.006
College or more	0.88	0.82	0.259

Fraction of participants who responded that they were satisfied with their care in followup survey presented by targeting dimension and treatment status. P-value for difference within targeting dimension, across treatment status presented in last column.

Table 10: How Cash Benefits Affect the Use of Informal Care

		Separately by state			
	Baseline	Arkansas	Florida	New Jersey	
Treatment group	117.91	114.23	125.29	117.25	
	(3.15)	(4.64)	(6.48)	(5.72)	
Control group	112.40	110.20	122.54	107.44	
	(3.37)	(4.94)	(7.01)	(6.10)	
Implied percent effect	4.90	3.66	2.24	9.13	
	(4.22)	(6.27)	(7.89)	(8.17)	
Observations	2,783	1,237	688	858	

The mean of informal care hours by treatment status are presented in the first two rows (standard errors in parentheses). The implied percentage impact of being in the treatment group reported in third row. Baseline column includes observations from all three states that conducted experiments. Observations gives the total (treatment plus control) number of observations for the column. * $p < .05$, ** $p < .01$

A The targeting benefit of in-kind transfers

This section presents a simple model that establishes two key results about the targeting effects of providing benefits in-kind relative to in cash. The first is that, other things equal, a means-tested cash-benefit program targets individuals with better health and with less costly informal care options—the opposite of the targeting that would be desired by a utilitarian planner. The second is that, other things equal, a means-tested in-kind benefit program that provides in-kind formal care (a substitute for informal care) targets individuals with worse health and with more costly informal care options, as desired by a utilitarian planner.

There are two periods. The individual has preferences over consumption in each period:

$$U = u(c_1) + u(c_2).$$

The rates of time preference and of interest are both zero. In the first period, the individual chooses how much to consume and save out of his (exogenous) after-tax resources:

$$c_1 + s_1 = (1 - \tau)\alpha,$$

where α is the individual's labor productivity, s_1 is first period savings, and τ is a labor tax. In the second period, the individual requires a fixed amount of care, η , which is known when the individual chooses how much to consume in the first period. The individual can meet this need by any combination of formal and informal care, which are perfect substitutes. The price of formal care is p . The (opportunity) cost of informal care is σ . In the second period, the individual consumes whatever resources are left over after receiving net government transfers, which may include both cash and in-kind formal care components, (t_c, t_f) , and paying for the (exogenous) level of care he requires, η . The second-period budget constraint depends on the exact nature of the government transfer programs, which are considered in turn now.

A.1 No benefit program, $(t_c, t_f) = (0, 0)$

Absent any government transfers in period two, the individual chooses whichever type of care is cheapest, i.e., he chooses formal care if and only if $p \leq \sigma$. The second-period budget constraint in this case is

$$c_2 = s_1 - \min\{\sigma, p\}\eta \equiv s_1 - \tilde{\sigma}\eta,$$

using $\tilde{\sigma} \equiv \min\{\sigma, p\}$. The individual optimally splits his resources equally between the first and second periods, so utility is

$$U(\alpha, \sigma, \eta; P = 0) = 2u \left(\frac{\alpha - \tilde{\sigma}\eta}{2} \right), \quad (3)$$

where $P = 0$ indicates that the individual does not participate in any government benefit program.

Next, we consider two types of schemes available to the planner given the information constraints she faces: a means-tested cash benefit program (similar to, for example, the US SSI program) and a means-tested in-kind benefit that provides formal home care, based on the US Medicaid program.

A.2 Means-tested cash benefit program

This program makes a cash transfer to the individual, the size of which is decreasing in the individual's savings: $t_c(s_1)$ with $t'_c(s_1) < 0$.²¹ Under this program, the individual's second-period budget constraint is

$$c_2 = s_1 + t_c(s_1) - \tilde{\sigma}\eta.$$

The simplest version of such a program sets a floor on net-of-transfer savings and implicitly taxes away any savings beyond that point among people who participate in the program:

$$t_c(s_1) = \max\{0, \bar{t}_c - s_1\}.$$

Under this version of the program, the individual's second-period budget constraint is

$$c_2 = \max\{\bar{t}_c, s_1\} - \tilde{\sigma}\eta.$$

Conditional on participating in the program, the individual should not save any resources into period two, since any savings simply serves to reduce transfers one-for-one and so does not increase second-period consumption. Thus, conditional on participating in the program, $c_1^*(\alpha, \sigma, \eta) = \alpha$ and $c_2^*(\alpha, \sigma, \eta) = \bar{t}_c - \tilde{\sigma}\eta$. The individual's indirect utility

²¹One could consider a related cash-transfer scheme whose means test is based on the individual's *net* assets after paying for any (formal) care rather than on the individual's gross assets before paying for care (as in the scheme under consideration). In fact, however, a scheme based on net assets after care expenses is equivalent to the in-kind formal care scheme considered next.

conditional on participating is

$$U(\alpha, \sigma, \eta; P = 1) \equiv u(\alpha) + u(\bar{B} - \tilde{\sigma}\eta)$$

The individual participates in the program if and only if

$$\begin{aligned} \Delta U^P &\equiv U(\alpha, \sigma, \eta; P = 1) - U(\alpha, \sigma, \eta; P = 0) > 0 \\ \iff u(\alpha) + u(\bar{B} - \tilde{\sigma}\eta) &> 2u \frac{\alpha - \tilde{\sigma}\eta}{2} \end{aligned} \quad (4)$$

The benefit of participating in the program is greater total lifetime consumption (by \bar{t}_c) from receiving the cash transfer. The cost of participating in the program is that the individual potentially (depending on his type) has to distort his consumption—consuming more in period one than in period two—in order to qualify.

The following types of people are more likely to participate in this program:²²

- Lower labor productivity α (good targeting)

It's more costly for high-labor productivity people to meet the means test, since they would otherwise wish to save more than lower-labor productivity people given their greater demand for c_2 .

- Lower care costs σ and lower care needs η (bad targeting)

It's more costly for high-care cost individuals to meet the means test (i.e., they would have to distort their consumption more), since they would otherwise wish to save more than people with lower care costs given their greater expenses in period two.

Like the first best, the means-tested cash benefit program redistributes toward those with lower labor productivity. Unlike the first best, however, this program redistributes away from those with higher care costs, which increases the risk to living standards from these costs. This occurs because the means test imposes greater costs on people with higher care costs, since in the absence of the program they would like to save more than otherwise similar people with lower care costs. This major limitation of means-tested cash benefit programs motivates the search for better alternatives. The commonly observed solution to this problem is a means-tested in-kind program, to which we now turn.

²²Note also that people with a greater willingness to substitute consumption over time are also more likely to participate in the program, since it is not as costly for them to have a non-smooth consumption profile. This is an example of “bad targeting”: having a high elasticity of intertemporal substitution, *ceteris paribus*, is not an indicator of having a high (lifetime) marginal utility of wealth. The first-best insurance program is neutral toward heterogeneity in preferences over intertemporal substitution.

A.3 Medicaid-type means-tested in-kind benefit program

As just seen, a major disadvantage of means-tested cash benefit programs is that they exacerbate the consumption risk arising from care cost risk. This raises the question of whether alternative program designs could increase insurance against care cost risk. One idea that may have potential in this regard is a means-tested in-kind program. One difficulty of paying cash benefits is that cash is as attractive to “mimics” as it is to intended beneficiaries, so cash benefits do not lead to any systematic self-selection into program participation of different types. Paying benefits in kind, on the other hand, has the potential to improve on this situation. Including in the benefits bundle goods (or even bads) that are valued more highly by intended beneficiaries than by potential mimics can create favorable self-selection into the program, i.e., selection that furthers the planner’s goal of redistributing to intended beneficiaries.

Consider a Medicaid-type program, which is means tested and provides benefits in kind, in the form of *formal* home health care. Specifically, suppose that Medicaid provides qualifying individuals up to μ units of formal care. In addition to paying benefits in kind, this program differs in one other important way from the cash-benefit program considered above: it tests means *net of (after)* spending on *formal* care, as opposed to gross of (before) spending on any care. Specifically, suppose that in order to qualify for Medicaid, someone’s wealth net of spending on formal care must not exceed \bar{w} .

If Medicaid benefits are sufficient to cover the individual’s care needs entirely, $\mu \geq \eta$,²³ then second-period consumption conditional on participating in Medicaid is:

$$c_2 = \min\{\bar{w}, s_1\}.$$

The incentives for saving facing individuals in this case are slightly different from those in the cash-benefit case. Now, saving up to \bar{w} is not implicitly taxed by the means test; individuals keep any savings they have up to \bar{w} for themselves. Saving beyond \bar{w} is taxed away completely by the means test. Utility conditional on participating is therefore

$$U(\alpha, \sigma, \eta; P = 1) = \begin{cases} 2u \frac{\alpha}{2} & , \quad \text{if } \alpha \leq 2\bar{w}; \\ u(\alpha - \bar{w}) + u(\bar{w}), & \text{if } \alpha \geq 2\bar{w}. \end{cases}$$

²³We focus on this case since it is much simpler than the case in which Medicaid benefits are not sufficient to cover the individual’s care needs, $\mu < \eta$, and since the results and intuition are virtually identical in the two cases.

The individual participates in the program if and only if

$$\begin{aligned} \Delta U^P &\equiv U(\alpha, \sigma, \eta; P = 1) - U(\alpha, \sigma, \eta; P = 0) > 0 \\ \iff U(\alpha, \sigma, \eta; P = 1) &= \begin{cases} 2u \frac{\alpha}{2}, & \text{if } \alpha \leq 2\bar{w}; \\ u(\alpha - \bar{w}) + u(\bar{w}), & \text{if } \alpha \geq 2\bar{w}. \end{cases} > 2u \frac{\alpha - \tilde{\sigma}\eta}{2} \end{aligned} \quad (5)$$

The following types of people are more likely to participate in this program:

- Lower labor productivity α (good targeting)

It's more costly for high-labor productivity people to meet the means test.

- Higher care costs σ and higher care needs η (good targeting)

The value of the Medicaid transfer is increasing in one's care costs, since participating in Medicaid eliminates the need to finance one's own care.

This program, unlike the cash-benefit program, transfers more resources to people with higher care costs and thus provides more insurance against this risk. This is the targeting benefit of paying benefits in kind. The improvement in targeting, however, does not come without a cost. The cost is that some people distort their care choices in order to increase the amount of benefits for which they qualify. This arises for two reasons. First, it arises because Medicaid provides benefits in the form of formal care. For people whose costs of informal care are lower than the cost of formal care, this is a social waste; society would be richer if these people received informal rather than formal care, but the incentives built into the in-kind benefits program lead them to consume more costly formal care instead. This is the consumption distortion that most people think of. The second reason Medicaid leads people to distort their care consumption is that its means test is based on net wealth after paying for *formal* care. This latter effect does not arise in the preceding analysis since it considers only individuals whose care needs are met completely by Medicaid's maximum benefit, μ . For individuals who require more care than μ , Medicaid's means testing rules provide an incentive to spend some of their own assets on formal care in order to qualify for Medicaid coverage (also of formal care) even if their costs of informal care are lower than the cost of formal care. In sum, by paying benefits in-kind, this program is better able to insure people against the risk of having high care costs but at the cost of encouraging people to consume too much formal care rather than informal care.

The cash and in-kind schemes also potentially differ in terms of the extent to which they distort the timing of consumption (and, more generally, any earlier decisions that affect future benefit receipt, such as labor supply). But, aside from being very difficult to

measure empirically, this distortion is common to all means-tested programs; it is not special to the issue of interest in this paper of whether to pay benefits in kind or in cash.