

# The demand for skills training among Medicaid home-based caregivers

Christopher J. Cronin\*  
University of Notre Dame

Ethan M.J. Lieber  
University of Notre Dame  
and NBER <sup>†</sup>

April 6, 2023

**ABSTRACT:** Medicaid spends nearly 100 billion dollars annually on home and community-based care for the disabled. Much of this care is provided by personal care aides, few of whom have received training related to the services they provide. We conducted a randomized controlled trial to estimate their demand for training. We find that 13 percent of these caregivers complete training without an incentive. Paying the caregivers four times their hourly wage increases training completion by roughly nine percentage points. Additional experimental variation suggests that among individuals confirmed to be aware of the training, the financial incentive increases completion from 35 to 58 percent. Demand curves based on these results suggest that while many caregivers value the opportunity to train, policies aimed at universal take up require large financial incentives.

**KEYWORDS:** Skills Training, Medicaid, Home and Community Based Services, Caregiver, Demand for Training, Marginal Value of Public Funds.

**JEL CLASSIFICATION:** D61, H00, I18, I21, J44, J88

---

\*Corresponding Author. Email: ccronin1@nd.edu. Address: Department of Economics, University of Notre Dame, 3053 Jenkins Nanovic Halls, Notre Dame, IN 46556.

<sup>†</sup>We thank Tamara Konetzka, Lee Lockwood, Corina Mommaerts, Hessel Oosterbeek, Meghan Skira, Yulya Truskinovsky, and Courtney Van Houtven as well as seminar participants at ASHEcon 2021, the University of Chicago, the 2021 National Tax Association Meetings, the University of Notre Dame, and Université Laval for helpful comments. Becca Brough, Connor Chapkis, and Brendan Perry provided excellent research assistance. We are grateful to the Wilson Sheehan Lab for Economic Opportunities for both research and financial support. This project would not have been possible without the cooperation of our partner organization, Helping Restore Ability, and numerous members of their staff, including Vicki Niedermayer, Girish Shelke, David Sabol, and Tiffany Denny. This study was registered with the AEA RCT Registry: AEARCTR-0007156.

# 1 Introduction

In 2018, Medicaid spent \$92b on home and community based services (HCBS), which is more than the US spent on the Supplemental Nutrition Assistance Program (“food stamps”), Temporary Assistance to Needy Families, the Earned Income Tax Credit, and numerous other social programs (Blom, 2019). This spending provides help to over 4 million disabled Americans annually (O’Malley Watts et al., 2020), a number that is sure to grow as the US population continues to age. HCBS funds largely support individuals who need assistance with day-to-day tasks like bathing, dressing, and preparing food, but prefer home-based care to living in a nursing home. There is concern that the quality of care provided by a large group of HCBS caregivers, personal care aides (“caregivers” hereafter), is sub-optimal since the vast majority receive no formal training related to the services they provide (LeBlanc et al., 2001; Spetz et al., 2019).<sup>1</sup>

In recent years, several policies have sought to promote caregiver training. The Personal and Home Care Aide State Training Demonstration Program, part of the Affordable Care Act, supported efforts to develop training for caregivers and led to free, online training programs in some states. President Biden’s Build Back Better agenda also called for investment in training programs to help caregivers deliver higher-quality care (H.R.5376). Despite the policy push to increase caregiver training, we know very little about the demand for training and consequently, how effective attempts to increase training will be.

To take first steps in this direction, we design, conduct, and evaluate a randomized controlled trial (RCT) that offers home-based caregivers the opportunity to complete an online training program. A random subset of participants are offered a \$50 incentive, equivalent to nearly four-times the average worker’s hourly rate, to complete the one hour training. This incentive effectively varies the price of the training and allows us to estimate a demand curve for the training. This demand curve is useful because it quantifies information on caregivers’ willingness to pay for training and provides a first step towards analyses of policies that encourage caregiver training. We find that among individuals not offered any financial incentive, 13 percent completed training. The \$50 incentive increased the completion rate by nearly 67 percent, or nine percentage points.

All of the caregivers in the experiment were contacted via email. In an effort to gauge the fraction that read the message and were therefore aware of the opportunity, we also offered a random subset of participants \$10 for simply replying to the email. Based on the replies we received and training completion rates, approximately 37 percent of caregivers were aware of the training opportunity. Among those who were aware of the training, we estimate

---

<sup>1</sup>There are multiple categories of both paid/formal and unpaid/informal home-caregivers. Moreover, Medicaid HCBS dollars fund a variety of care categories, including personal care aides, home health caregivers, hospice care, adult day care, etc. Training requirements are largely nonexistent for personal care aides, as well as informal caregivers, though training is commonly required in the other care categories. See Section 2 for more information.

that between 35 and 51 percent of caregivers not receiving a financial incentive completed training, while 59 to 65 percent of incentivized caregivers completed. These results suggest that for training programs like the one used in our experiment, a policy that makes training freely available, and potentially offers a financial incentive for completion, can result in a majority of caregivers taking up the training without great expense. However, if the policy aims to have 100 percent participation, a large financial incentive is required.

We conclude our analysis by using the results from our experiment and the Marginal Value of Public Funds (MVPF) framework (Finkelstein and Hendren, 2020) to evaluate the welfare implications of a policy that offers caregivers a free online training program.<sup>2</sup> Such a policy may be justified in this setting by an externality that could prevent efficient levels of training: Many HCBS recipients are insured through Medicaid, meaning at least part of the benefit of training—higher quality care, which results in lower health care costs—is captured by the government insurer and not the caregiver or care recipient.<sup>3</sup> We draw two main conclusions from our analysis. First, the policy is likely to increase welfare per dollar spent more than increasing caregiver wages in Medicaid, the primary payer for home-based caregiving. Key to this finding is the fact that offering online training is relatively inexpensive (the cost-equivalent increase in Medicaid caregiver wages is less than two pennies per hour) and our experiment suggests that many caregivers will choose to complete training even without a financial incentive. Second, if training caregivers reduces healthcare expenditures even by a small amount, then the policy is likely to cost less to finance than the healthcare expenditures saved. As with any welfare analysis, these conclusions require a number of strong assumptions, which we highlight and discuss below.

Our work contributes to literatures in economics, health policy, and medicine studying various aspects of the long-term care market including insurance (e.g., Brown and Finkelstein, 2011; Mommaerts, 2020), hospital incentives (e.g., Einav et al., 2018), the supply of caregivers (e.g., Spetz et al., 2015, 2019), and trade-offs associated with informal care provision (e.g., Van Houtven et al., 2013; Skira, 2015; Barczyk and Kredler, 2018; Mommaerts and Truskinovsky, 2020). Although skills training has been shown to improve knowledge, skills, and patient outcomes for physicians and nurses (Davis et al., 1995; Marinopoulos et al., 2007; Khatony et al., 2009), the strand of this literature focusing on caregiver training is quite limited. Cooper et al. (2017) conduct a meta-analysis of papers that evaluate the

---

<sup>2</sup>The MVPF of a particular policy measures the beneficiaries’ willingness to pay for the policy divided by the net cost to the government; thus, it measures the benefit per dollar spent. Because the measure does not internalize market distortions stemming from raising government revenue, it is most useful when comparing two policies. The policy with the larger MVPF generates more welfare for marginal changes in the programs’ budgets.

<sup>3</sup>The theoretical literature on skills training has studied other impediments to efficient investment in the form of worker liquidity constraints (Becker, 1962), imperfectly competitive labor markets (Stevens, 1994; Acemoglu and Pischke, 1999), and various forms of asymmetric information (Chang and Wang, 1996; Acemoglu and Pischke, 1998).

efficacy of training programs for paid home caregivers (broadly defined) and find just six quantitative papers, only one of which they deem high quality.<sup>4</sup> More recently, Van Houtven et al. (2019) conducted an RCT that provided training to informal, family caregivers of military veterans and found that training increased experienced quality of care. The authors also examined veterans’ subsequent use of health care along a number of dimensions, but lacked the statistical power to reject the null for fairly large point estimates. For example, veterans whose family caregivers received training were approximately 10 percent less likely to have an emergency department visit or a hospitalization, but these effects were not statistically different from zero. Our RCT complements this past work by providing estimates of the demand for training in the context of a newly available, relatively short online training module.

We also contribute to a strand of literature on the supply of and demand for skills training (see Leuven and Oosterbeek, 1997, for a review). Demand estimation in our paper is most similar to Hidalgo et al. (2014), which randomized vouchers for training programs to a set of low-skilled individuals in the Netherlands. The authors found that lowering the price of training by approximately €21 per hour (\$30 in real 2021 U.S. dollars) increased training completion by 44 percent.<sup>5</sup> We find broadly similar results, a \$50 price reduction increased training completion by 67 percent. Despite the large differences in contexts and the content of the training, this similarity suggests that our results may be more broadly informative than otherwise expected.

The remainder of this paper is organized as follows: Section 2 provides background information on caregiving in the US. Section 3 describes our partner organization and the experiment. Section 4 discusses our data and empirical strategy. We present our findings in Section 5 and discuss the implied demand curve in Section 6. Our welfare analysis can be found in Section 7. Finally, we discuss our findings and summarize the limitations of our analysis in Section 8.

## 2 Caregiving and Caregiver Training

Medicaid is the primary payer for long-term care in the United States and spending on long-term care has risen considerably over time (O’Malley Watts et al., 2020). As seen in Figure 1,

---

<sup>4</sup>The single high quality paper (King et al., 2012), features an RCT that randomly assigns patients to a package of benefits that includes training for their caregiver, but also increases caregiver-patient contacts, supervision of caregivers by medical professionals, and an activity of daily living evaluation and exercises to facilitate improvements. The study found significant improvements in the overall health and mental health (SF36 measures) of treated patients 7 months post treatment, but these improvements obviously cannot be attributable to training alone. Cooper et al. (2017) find three papers examining training programs that do not include follow-up supervision. Caregiver outcomes such as topical knowledge, job satisfaction, and resilience varied across the studies. Patient outcomes were not studied.

<sup>5</sup>We thank Hessel Oosterbeek for generously providing additional statistics necessary to make this calculation.

Medicaid’s long-term care spending increased from an inflation adjusted \$71 billion in 1995 to \$146 billion by 2016. Historically, the majority of these dollars went towards nursing home care. Beginning in the 1990s, there was a movement away from placing disabled individuals in nursing homes towards keeping them in their homes where supports and services could be provided at a (potentially) lower cost. Since 2013, these home and community-based services (HCBS) constitute the majority of Medicaid programs’ long-term care dollars. In 2018, the most recent data available, Medicaid spent 57 percent, approximately \$92 billion, of its long-term care dollars on HCBS for 4 million enrollees (O’Malley Watts et al., 2020).




Figure 1: Medicaid Long-term Care Expenditures

Notes: Authors’ calculations based on data in Eiken et al. (2018). Expenditures are in 2010 dollars.

A large share of Medicaid HCBS funds go to home care workers, who fall into two primary groups: (i) personal care aides (PCAs) and (ii) home health aides (PHI, 2016). A third group, nursing assistants, perform tasks similar to home health aides, but make up a small fraction of home care workers. PCAs provide help with activity of daily living (ADL) limitations such as dressing, bathing, or toileting, as well as instrumental activity of daily living (IADL) limitations such as shopping, cooking, or any of the multitude of activities that individuals carry out while living independently. Importantly, these PCAs do not provide medical care. Home health aides do provide medical care that could normally be done in a doctor’s office (e.g., dressing wounds, administering shots, etc.), but may also help with ADLs and IADLs.

A large share of Medicaid’s \$92b annual spend on HCBS likely goes to PCAs, but CMS

does not supply the exact figures. Surprisingly, a 2017 report by the U.S. Government Accountability Office could not determine Medicaid HCBS expenditures on (nor recipients of) PCA services alone (Iritani, 2017). O’Malley Watts et al. (2020) report that 1.2 million Medicaid HCBS enrollees received personal care services through a state plan, but the true number is likely much larger than this as (i) this number includes reports from just 35 states and (ii) in all states, a large share of HCBS recipients are enrolled in waivers, which also fund PCA services. Using IPUMS data and occupational codes suggested by Osterman (2017), we estimate that there are three PCAs for every two home health aides in the population at large in both 2016 and 2017, but these figures may differ in the Medicaid HCBS population.<sup>6</sup>

To gain access to Medicaid HCBS in Texas—where our experiment was conducted—a person first needs to be put on a program-specific waiting list. As of 2017, over 280,000 Texans were on a waiting list, which is the largest number among US states by far (Musumeci et al., 2019).<sup>7</sup> Once an individual reaches the top of a list, Medicaid determines whether the individual is eligible for services and for which types. Eligible individuals then meet with a Medicaid official to choose a home care agency, which administers benefits.

Individuals receiving PCA services through Medicaid HCBS obtain a caregiver by choosing one of two pathways. In the first, the home care agency provides the HCBS recipient with a caregiver that is employed by the home care agency. This is called the “Personal Care Attendant Services” program. In the second, the HCBS recipient chooses/hires the caregiver herself, maintaining the official role of employer. We refer to this as the “Consumer Directed Services” (CDS) program.<sup>8</sup> An advantage of the CDS program is that it allows HCBS recipients to hire friends or family members as caregivers, many of whom have been providing care since long before receiving Medicaid funds. Though the HCBS recipient is technically an employer in this situation, Texas Medicaid stipulates that the HCBS recipient must hire a home care agency to manage employer-specific administrative tasks, such as payroll, taxes, bookkeeping, etc.

While home health aides must be trained and certified (Osterman, 2017), the PCAs that our experiment focuses on tend to have little formal training related to the services they provide. There are no federal training or certification requirements for PCAs (Spetz et al., 2019). As reported in LeBlanc et al. (2001), no more than thirteen percent of states required

---

<sup>6</sup>We use the IPUMS occupation code 4610 for personal care aides, 3600 for home health aides, and require workers of both types to work in IPUMS industry codes 7580, 8170, 8180, 8370, 9290, and 9480. It is important to note that many informal caregivers are not represented in this calculation because caregiving may not be their listed occupation.

<sup>7</sup>Discussions with our RCT partner and reports in the popular press (e.g. Eiserer, 2019) suggest individuals spend an average of seven years on the wait lists and can spend considerably longer for particular programs.

<sup>8</sup>The CDS program is actually a collection of multiple programs targeted at specific subpopulations (e.g., in Texas, the Texas Community Based Alternatives waiver, the Community Living Assistance & Supportive Services waiver, etc.). All of the programs share the same general feature that the HCBS recipient is the employer. For ease of exposition, we will refer to the collection of policies as the CDS program.

all home care workers to undergo some form of training.<sup>9</sup> More recently, a survey of state regulations found that CDS personal care aides were not required to have any training in thirty-eight of the fifty states (PHI, 2021). Among the twelve requiring training, none require training for private-pay arrangements. Training is only required if services are paid for with Medicaid funds. We will show below that many of the PCAs in our estimation sample were previously unpaid, family caregivers and Burgdorf et al. (2019) find that less than 8 percent of these caregivers receive role-related training.

Because of the strong relationships between many caregivers and care recipients, we might expect caregivers to seek training on their own. On the other hand, caregivers are not able to capture all of the benefits that arise from training and, as a consequence, may invest less than the socially optimal amount in training. For example, consider improvements in a patient’s health that could result from a caregiver’s training. This leads to less time at the doctor’s office for the patient and lower health care expenditures. The patient obtains the benefit of less time in the doctor’s office and, in some models, could transfer that utility to the caregiver. However, because the patient is likely paying only a small fraction of her health care costs, the insurer likely captures much of the benefit from reduced health care expenditures. Since the reduction in expenditures is not paid to the caregiver, there is an externality that leads to underinvestment in training in the private market and provides a rationale for policies which subsidize training.

## 3 Description of Experiment

### 3.1 Partner Organization and Population

Our partner organization is Helping Restore Ability (HRA), a home care agency located in North Texas. HRA assists Medicaid HCBS recipients by either (i) providing the recipient with a caregiver that HRA employs or (ii) acting as a financial services management agency for a caregiver that the recipient employs (i.e., CDS program). All caregivers are personal care aides. In Texas, a home care agency must have a training program in place for the caregivers it employs to maintain a license with Medicaid. HRA requires the caregivers it employs to complete approximately 12 hours of online training each year. These training modules cover topics such as infection prevention and control and recognizing elder abuse. Caregivers employed through a CDS program are neither required nor offered the opportunity to complete these same trainings. For expositional ease, we will refer to both HRA’s directly employed caregivers and CDS caregivers collectively as HRA’s caregivers.

---

<sup>9</sup>The term “home care workers” refers to more than just PCAs. It also includes home health care aides, certified nursing assistants, physical therapists, and others. However, PCAs are the least skilled home care workers and tend to have the lowest requirements (PHI, 2021). This suggests that the requirement that *all* home care workers have some form of training likely reflects training requirements for PCAs.

Table 1 highlights several demographic and employment characteristics of the caregivers in the experiment. A large majority, 89 percent, are working through the CDS program rather than as HRA employees. Eighty percent are female, 28 percent are Black, and 19 percent are Hispanic. Ninety-two percent of caregivers serve just one patient; however, most patients have multiple caregivers, meaning just 45 percent of caregivers share no common patients with other caregivers. The average caregiver is 44 years old. The average hourly wage is \$12.16.<sup>10</sup>

Table 1: Summary Statistics

	Mean (1)	Standard Deviation (2)
<i>Binary variables</i>		
Consumer directed services	0.89	0.31
Female	0.80	0.40
Black	0.28	0.45
Hispanic	0.19	0.39
Family/friend	0.69	0.46
One patient	0.92	0.27
One caregiver	0.45	0.50
<i>Continuous variables</i>		
Age (years)	43.71	14.90
Months with patient	85.35	95.52
Months with HRA	41.42	44.40
Hours per week	30.32	17.68
Wage rate	12.16	3.91

Notes: Means and standard deviations provided for demographic variables. Each of the variables at the top of the table are binary; variables in the lower portion are continuous. The distributions of months with the patient and months with HRA have a long right tail. The medians for those variables are smaller than the means, 54 and 26 months respectively.

Sixty-nine percent of HRA’s caregivers are family members or were friends of the patient prior to providing care, which is consistent with the observation above that most caregivers are working through the CDS program. The average patient-caregiver relationship pre-dates the caregiver’s relationship with HRA by 44 months, suggesting that many caregivers likely provided informal care to the patient prior to qualifying for Medicaid HCBS funds. Given this information, we might expect the caregivers to be similar to the informal caregiver population. Based on summary statistics from Mommaerts and Truskinovsky (2020), our caregivers are similar to informal caregivers on some dimensions, but are more likely to be

<sup>10</sup>The Medicaid waiver pays a fixed amount per patient based on need. CDS patients have some ability to negotiate wages with their caregivers, though the state caps wages within each specific waiver program. Patients can “afford” to pay higher hourly wages by reducing hours of care.



female, less likely to be white, and provide many more hours of care on average. Demographically, our sample appears to be similar to the general population of formal caregivers, which we measure using the IPUMS coding of the American Community Survey (Ruggles et al., 2021). See Appendix A for details. It is not entirely surprising that our caregivers are similar demographically to formal caregivers since consumer-directed programs have become common in Medicaid HCBS. For example, of the 34 states that offer a personal care state plan option, 20 of them allow consumer-direction of services (O’Malley Watts et al., 2020). Though formal statistics are hard to find, consumer-directed programs appear to be common among HCBS services funded through waiver programs as well.

## 3.2 Experiment

The experiment directed caregivers to a training module on the website RCTCLEARN.NET. RCTCLEARN.NET began as a partnership between healthcare agencies and academic institutions in Minnesota, but since 2006, has been part of a joint venture with the National Association for Home Care and Hospice to provide internet-based training to home care and hospice workers. RCTCLEARN.NET creates training modules and provides access to the content, as well as other services like tracking trainee progress, to home care associations and home care agencies for a fee. HRA has access to the RCTCLEARN.NET platform through its membership in the Texas Association for Home Care and Hospice, which has purchased access to RCTCLEARN.NET’s materials.

The module we use in the experiment, *Effective Communication Using the SBAR Tool*, focuses on helping the caregiver communicate with the care recipient and medical workers. This training module could reduce coordination costs, which are high in healthcare (Cebul et al., 2008), and potentially lead to better patient outcomes.<sup>11</sup> The module is set up similar to many human subjects or CITI trainings: there are (web)pages of information presented to a reader with a quiz at the end. If the trainee gets at least 80 percent of the questions on the quiz correct, she passes and is offered a printable certificate. To our knowledge, these certificates are not required by any home care agencies when applying for employment and do not qualify employees for raises or other awards.

We assigned 747 caregivers to four groups, which are depicted in Figure 2. Our control group was sent an email with the information necessary to complete the training module. Our three treatment groups received very slightly altered versions of that same email. Treatment Group 1 was offered a \$50 Visa gift card for passing the quiz and a \$10 Visa gift card for responding to the email. Treatment Group 2 was offered a \$50 Visa gift card for passing the quiz. Treatment Group 3 was offered a \$10 Visa gift card for responding to the email. The

---

<sup>11</sup>*SBAR* stands for Situation-Background-Assessment-Recommendation. Among the modules not already completed by the caregivers HRA employs, HRA preferred this module because they thought it stood the best chance to improve the patient and caregiver relationship.

exact text of the emails for each group can be found in Appendix B. While the reasoning behind offering individuals a \$50 incentive to complete the training is fairly straightforward, the \$10 incentive to reply to the email may be less so. The \$10 incentive, located at the very end of the emails sent to treatment groups 1 and 3, is designed to help us estimate what fraction of individuals are aware of the training opportunity.

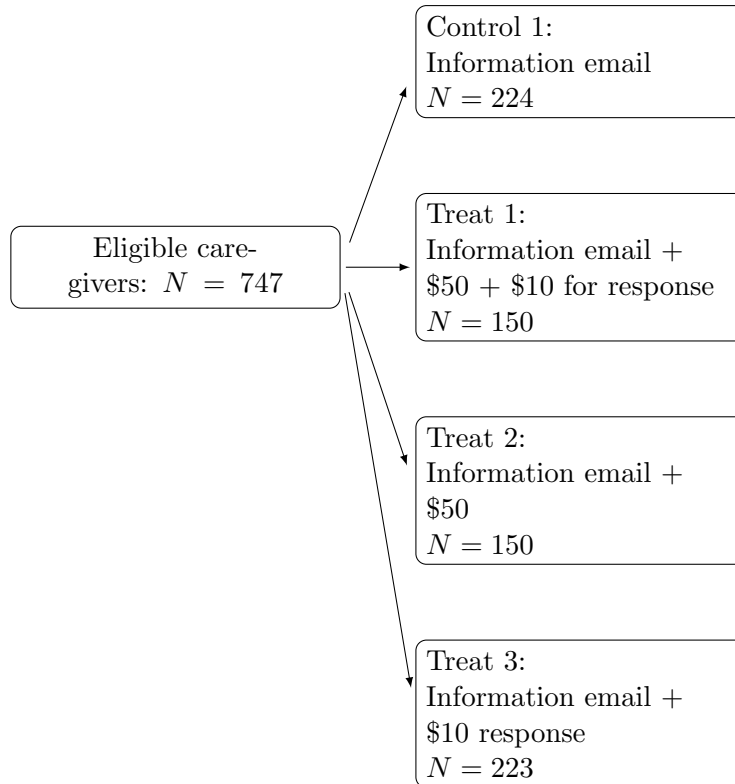


Figure 2: Experimental Groups

Emails were sent to participants each Monday, Thursday, and Saturday between February 15th and March 7th, 2021.<sup>12</sup> Because caregivers do not work standard 8–5 schedules, we alternated the time at which emails were sent between 8 AM and 4 PM. Upon completing training, trainees were removed from the list-serve. Upon replying to our email, the \$10 email reply incentive language was removed for individuals in treatment groups 1 and 3.

Following Banerjee et al. (2020), we used rerandomization to assign participants to treatment and control groups. This process increases the likelihood that the treatment and control groups are balanced on observables (and unobservables) while still reaping the benefits of randomized assignment. To be explicit, we randomly assign individuals to treatment and control groups and test for balance on a set of  $K$  variables. We follow the suggestion of Banerjee et al. (2020) that the number of randomizations be  $\min\{100, N\}$  where  $N$  is

<sup>12</sup>Our pre-analysis plan specified a two-week time frame for the experiment from February 15th - February 28th. We deviated from that plan and extended the training opportunity for one week because of the snow and ice storm which knocked out power to large swaths of Texas, including many of the participants in our experiment. This is discussed further in Section 5 where we show it does not materially affect our results.

the number of units in the experiment; thus, we randomize 100 times and select the randomization with the lowest maximum t-statistic. We chose a broad set of variables for the rerandomization procedure that are meant to capture caregiver characteristics that could be correlated with propensity to take up the training. These include variables for (i) whether the caregiver was part of a CDS program; (ii) age, gender, race, ethnicity; (iii) if the caregiver is a family member or friend of the patient; (iv) months the caregiver has been helping the patient and working with HRA; (v) if the caregiver was actively providing care; (vi) if the caregiver’s patient was insured through Amerigroup; (vii) if the caregiver had completed a training module on RCTCLEARN.NET as part of a previous experiment; (viii) indicators for whether the caregiver had responded to the online, mail, or phone survey when providing demographic information; and indicators for whether variables (ii)-(iv) had missing values.

Several variables require explanation. The eight percent of caregivers that have patients insured through Amerigroup are isolated because in the four years preceding our experiment, Amerigroup offered some of these caregivers financial incentives to complete training and used HRA to administer the program. As such, these caregivers have experience with both training and facing financial incentives for participation, which could impact completion rates. Similarly, we conducted a separate experiment with HRA caregivers in the Fall of 2020, which is described in Appendix C. Like the caregivers HRA employs directly and Amerigroup caregivers, those completing training as part of this earlier experiment have more experience with both training and the RCTCLEARN.NET platform, which could impact training here.<sup>13</sup> Caregivers are coded as “actively providing care” if HRA records indicate that the caregiver had received a pay check in either of the previous two pay cycles preceding the fall experiment. Due to both COVID-19 and some patients requiring seasonal care, 40 percent of caregivers were not active at the time of the experiment. Finally, HRA issued a survey to all caregivers preceding the experiment to supplement the demographic information they already had on file. We view the medium used by caregivers to complete the survey as a measure of their comfort with internet technologies, which could impact completion rates.

Table 2 shows the balance tests between the control group and the various treatment groups. The first column provides the mean value of the variables used in the rerandomization procedure for the control group. The following three columns present the p-value of the differences in means between the specified treatment group and the control group. None of the p-values indicate statistically significant imbalances between the treatment and control groups at the five percent level. Wage rates might be slightly different across some groups, so we conduct our analyses with and without variables not included in the rerandomization procedure (such as wages) and see that the results are not materially affected.

This study was approved by the University of Notre Dame’s Institutional Review Board

---

<sup>13</sup>In additional analyses, we find no evidence that treatment effects vary by the previous experiment (see Appendix C), whether the caregiver was employed directly by HRA, or Amerigroup status (see Section 5.1).

(Protocol: 20-04-6024). We registered the study in the AEA RCT Registry (AEARCTR-0007156) on February 10, 2021, prior to randomization and launch. Additional details are provided in Appendix D.

Table 2: Balance Among Treatment and Control Groups

		P-value of Difference		
	Control Mean	Treatment 1	Treatment 2	Treatment 3
	(1)	(2)	(3)	(4)
<i>Panel A: Rerandomization variables</i>				
Consumer directed services	0.91	0.894	0.728	0.130
Age (years)	43.16	0.742	0.493	0.581
Female	0.76	0.211	0.701	0.172
Black	0.31	0.754	0.643	0.139
Hispanic	0.21	0.690	0.976	0.258
Family/friend	0.67	0.321	0.496	0.427
Months with patient	89.68	0.353	0.837	0.645
Months with HRA	42.02	0.891	0.702	0.470
Providing care	0.60	0.837	0.642	0.151
Amerigroup patient	0.07	0.860	0.758	0.137
Previous online module	0.24	0.864	0.748	0.933
Online survey	0.33	0.143	0.881	0.528
Mail survey	0.23	0.862	0.390	0.434
Telephone survey	0.14	0.526	0.795	0.331
Missing age	0.04	0.379	0.628	0.129
Missing female	0.38	0.389	0.260	0.875
Missing race	0.29	0.131	0.621	0.976
Missing ethnicity	0.29	0.131	0.621	0.976
Missing family/friend	0.34	0.196	0.308	0.869
Missing months with patient	0.34	0.602	0.511	0.816
Missing months with HRA	0.07	0.351	0.813	0.579
<i>Panel B: Additional variables</i>				
One caregiver	0.46	0.869	0.676	0.743
One patient	0.92	0.521	0.950	0.608
Female patient	0.43	0.971	0.270	0.892
Hours per week	31.98	0.223	0.424	0.361
Wage rate	12.63	0.070	0.717	0.083
Missing wage	0.42	0.697	0.526	0.157
Missing hours	0.42	0.697	0.526	0.157

Notes: Column (1) presents the mean of each variable for the control group. Columns (2)-(4) contain the p-value of the test that the mean for the control group is the same as the mean for each treatment group. Treatment 1 included the offer of \$50 for completing the module and the offer of \$10 for responding to the email; treatment 2 included the \$50 module completion offer; treatment 3 included the \$10 response offer.

## 4 Data and Empirical Strategy

Caregiver email addresses and demographic data were provided by HRA. Our outcome measures, whether the caregiver passed the quiz with a score of 80 percent or greater and whether the caregiver began the training module, come from RCTCLEARN.NET.

Our empirical strategy is a simple linear regression of the form

$$y_i = T_i\beta + X_i\Gamma + \varepsilon_i \quad (1)$$

where  $y_i$  is one of our two outcomes,  $T_i$  is a vector of dummies indicating which treatment was received (control group omitted),  $X_i$  is a vector of controls, and  $\varepsilon_i$  is the remaining error term. Initially, only the rerandomization variables are included in  $X_i$ ; subsequent regressions add controls for the patient’s gender, whether the patient had only one caregiver, whether the caregiver had only one patient, the caregiver’s hourly wage and hours worked, and controls for whether those variables are missing. Bruhn and McKenzie (2009) show that the rerandomization procedure affects inference and that including the rerandomization variables as controls leads to appropriately sized statistical tests. We use Eicker-Huber-White robust standard errors.

Because we are interested in estimating how the monetary incentives affect behavior, we will primarily focus on the following specification,

$$y_i = \alpha_0 + \alpha_1 T_{50} + \alpha_2 T_{10} + X_i\Gamma_0 + \nu_i. \quad (2)$$

Everything in Equation (2) is the same as in Equation (1) except that  $T_{50}$  is an indicator for whether or not the individual was offered \$50 to pass the quiz at the end of the training (treatment groups 1 and 2) and  $T_{10}$  is an indicator for whether or not the individual was offered \$10 to respond to the email (treatment groups 2 and 3). This regression specification directly estimates how subsidizing training (lowering the price by \$50) changes the participation rate ( $\alpha_1$ ).

## 5 Results

Table 3 presents the regression results from Equation (1). Columns (1)-(2) show estimates where the dependent variable is an indicator for whether or not the caregiver began the training (regardless of whether training was completed); columns (3)-(4) show estimates where the dependent variable is whether the training was completed with a passing grade on the quiz. Seventeen percent of the control group began the training module and 12.9 percent completed it. Focusing first on columns (1) and (3), we find that being randomized to the group that was offered the \$50 and \$10 incentives increased beginning training by 9.1 percentage points and training completion by 9.6 percentage points. Those point estimates translate into very large percent increases, 54 percent and 74 percent respectively. Being randomized to the group that was offered just the \$50 incentive to complete the training increased beginning and completing the module by 11.6 percentage points. Being randomized to the group that was offered just the \$10 incentive to respond to the email led to a statisti-

cally insignificant 3.9 percentage point increase in beginning the training and 4.2 percentage point increase in completing the training. Adding in additional controls for patient gender, patient/caregiver network, and caregiver hours worked and hourly wages (columns (2) and (4)) does not materially affect the results.

Table 3: Impacts of Treatments on Beginning and Completing Training

	Began training		Completed training	
	Rerandomization Controls (1)	Additional Controls (2)	Rerandomization Controls (3)	Additional Controls (4)
Treat 1: \$50 + \$10	0.091 (0.036) [0.012]	0.094 (0.036) [0.009]	0.096 (0.034) [0.005]	0.098 (0.034) [0.005]
Treat 2: \$50	0.116 (0.041) [0.005]	0.116 (0.041) [0.005]	0.116 (0.038) [0.002]	0.116 (0.038) [0.003]
Treat 3: control + \$10	0.039 (0.034) [0.262]	0.035 (0.034) [0.307]	0.042 (0.031) [0.182]	0.040 (0.032) [0.208]
Control mean	0.170	0.170	0.129	0.129
R-squared	0.263	0.273	0.279	0.282
Observations	747	747	747	747

Notes: Dependent variable is whether caregiver began training (columns (1) and (2)) or completed training (columns (3) and (4)). Columns (1) and (3) only include the rerandomization variables as controls. Columns (2) and (4) add controls for the patient’s gender, whether the patient had only one caregiver, whether the caregiver had only one patient, the caregiver’s hourly wage and hours worked, and controls for whether those variables are missing. Robust standard errors are shown in parentheses; p-values are presented in brackets.

The differences in estimated impacts between the two outcome variables are slight. More than 81 percent of individuals who began the training module successfully completed it and treatment status does not appear to have large impacts on whether or not an individual completes the training (conditional on beginning it). More precisely, if we run the same regression reported in column (4) of Table 3 but condition on having begun the training, none of the treatment indicators are statistically distinguishable from zero. As such, for the duration of the paper we will only include results in the main text for whether or not the person completes training and will relegate corresponding results on beginning training to Appendix E.

Table 4 presents results from estimating Equation (2). When only the rerandomization controls are included, we find that being offered \$50 to complete the training increased participation by 8.5 percentage points. When the additional controls are included, the estimated effect rises slightly to 8.7 percentage points. Relative to the control group’s completion rate

of 12.9 percent, offering caregivers \$50 to complete the training increases participation by 67 percent. The impact of the \$10 response incentive on completion is relatively small and statistically indistinguishable from zero, which provides some comfort that the inclusion of this treatment did not have large behavioral effects on participants (e.g., completing training out of guilt).

Table 4: Estimating the Impact of the \$50 Offer

	Rerandomization Controls	Additional Controls
	(1)	(2)
\$50 to complete	0.085 (0.027) [0.001]	0.087 (0.027) [0.001]
\$10 to respond	0.017 (0.026) [0.516]	0.017 (0.026) [0.523]
Control mean	0.129	0.129
R-squared	0.277	0.281
Observations	747	747

Notes: Dependent variable is whether caregiver completed training or not. Column (1) only includes the variables used in the rerandomization as controls. Column (2) adds controls for the patient’s gender, whether the patient had only one caregiver, whether the caregiver had only one patient, the caregiver’s hourly wage and hours worked, and controls for whether those variables are missing. Robust standard errors are shown in parentheses; p-values are presented in brackets.

In Appendix Table D2 we explore heterogeneity in the effect of the \$50 incentive on completion across a range of observables; however, this analysis is under powered.

## 5.1 Robustness

Our main effects are robust to a number of concerns. First, there is potential for spillovers across caregivers in the experiment. As seen in the summary statistics, approximately half of the Medicaid HCBS recipients in our sample had multiple caregivers. Because of this, it is possible that the caregivers for a single individual would communicate information about the experiment to the patient or simply talk between themselves about the experiment. This could either lead to greater participation (e.g. one caregiver helps the other complete the training) or lower participation (e.g. a control group caregiver becomes upset upon learning that another caregiver was offered \$50 to complete the training and so becomes less likely to complete the training herself). We discussed this possibility with HRA prior to the experiment. They did not believe this would be a problem because in their experience,

the caregivers rarely communicate with each other and do not have overlapping times at the patient’s home. We can more formally evaluate this possibility in a regression where we restrict the sample to caregivers whose patient only had a single caregiver—a subsample that is unlikely to be affected by the types of spillovers described. We present results from this exercise in the second column of Table 5. The results are very similar to our overall results, suggesting that this type of spillover is not playing a large role in our results.

Another concern is that caregivers directly employed by HRA may be affected differently from caregivers employed via the CDS option. HRA employees might be more likely to see the offer of training as an implied request that they complete the training, likely reducing the impact of the monetary incentive. At the same time, because HRA requires these caregivers to complete online training modules each year, this group is likely more familiar with this type of training. This could lead the costs of completing the training to be lower and potentially alter the impact of the \$50 incentive. Ex ante, it is not obvious in which direction these effects might push the estimate relative to CDS employees. In the third column of Table 5, we report results in which we restrict the sample to CDS caregivers. The results are very similar to our baseline results, suggesting the directly employed caregivers react to the incentives similarly to the CDS caregivers.

Caregivers for the 63 patients insured through Amerigroup (8 percent of our sample) had been asked to participate in an online training program in the past.<sup>14</sup> If experience with online training lowers the cost or affects the perceived value of completing the training, then this group of caregivers might react to the \$50 incentive systematically differently. As seen in the fourth column of Table 5, excluding the Amerigroup individuals does not materially affect the results.

At the start of our experiment, an ice and snow storm knocked out power to much of Texas. Because a large part of our study population was likely to be without running water, heat, or internet for a little less than a week, we extended the original deadline for the study by a week. However, we can observe the dates on which individuals completed the training module and so we are able to construct a version of our outcome variable based on the original time frame of the experiment. In column (5) of Table 5, we present results for a dependent variable which indicates the caregiver completed the training within the original two week time period. The fraction of caregivers in the control group who completed the training falls slightly to 12.1 percent, but the estimated impact of the \$50 monetary incentive is very similar to our original estimate. Thus it does not appear to be the case that lengthening the study time period had any meaningful impact on the results.

We have also estimated our main regressions as probits to check whether our assumed

---

<sup>14</sup>Texas’s Medicaid program is administered through a managed care system in which private insurers are paid by the state to provide coverage to Medicaid recipients. The previous training was part of a value-based partnership between Amerigroup and HRA.



linear functional form is driving the results. Both the linear and probit results (the average marginal effect in the sample is reported for the probit regressions) are reported in Appendix Table E2. The marginal effects for the two methods are quite similar, suggesting that the particular functional form assumed is not having dramatic effects on the estimates.

Table 5: Robustness of Main Results

	Baseline	Single Caregiver	CDS Only	Not Amerigroup	Two Week End Date
	(1)	(2)	(3)	(4)	(5)
\$50 to complete	0.087 (0.027) [0.001]	0.092 (0.044) [0.037]	0.086 (0.028) [0.002]	0.088 (0.028) [0.002]	0.086 (0.026) [0.001]
Control mean	0.129	0.157	0.127	0.120	0.121
R-squared	0.281	0.304	0.235	0.266	0.271
Observations	747	339	668	684	747

Notes: Dependent variable is whether caregiver completed training or not. The regression specification is the same as that used in column (2) of Table 4 except as noted in column headings. Column (1) repeats the results from column (2) in Table 4. Column (2) restricts the sample to patients with a single caregiver. Column (3) restricts the sample to caregivers in CDS programs. Column (4) excludes caregivers of patients insured through Amerigroup. Column (5) changes the dependent variable to whether the caregiver completed training within the first two weeks of the experiment. Robust standard errors are shown in parentheses; p-values are presented in brackets.

## 6 Demand for Training

In this section, we use the results from our experiment to estimate the training demand curve. We can directly measure the fraction of individuals who take up training at the randomized prices of zero and  $-50$  (the control and treatment groups respectively). As seen in the previous section, these fractions are 0.13 and 0.22. Assuming demand is linear, we can trace out the demand curve from these two points as seen in Figure 3, labeled ‘Initial’.

This demand curve implicitly assumes that all of the caregivers were aware of the training opportunity. However, this seems unlikely since we contacted caregivers via email: Analysis by Mailbird, an email management software company, suggests that 35 percent of emails are never read.<sup>15</sup> If everyone were aware of the opportunity, it seems likely that a larger fraction of individuals would complete the training at any given price and so our realized demand curve would potentially rotate and/or shift up. This is especially important when we are thinking about the external validity of our results.

In light of expected incomplete awareness, we included in our experimental design a low-cost method by which caregivers could signal program awareness. Half of the caregivers in

<sup>15</sup><https://www.getmailbird.com/email-use-statistics/>





Figure 3: Estimated Demand Curves

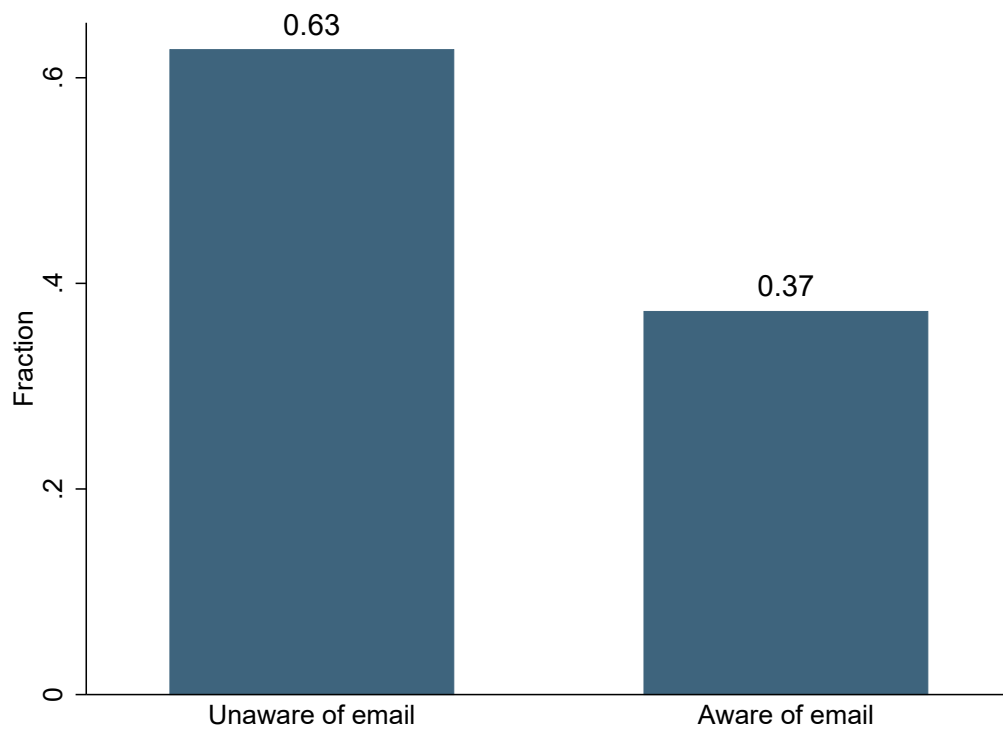
our study (i.e., treatment groups 1 and 3) were randomly selected to receive an email that included text at the bottom that reads, “In addition to providing training information, we are trying to confirm that our emails are being received by caregivers. If you reply to this email with ‘got it,’ we will send you a \$10 Visa gift card.” Figure 4a shows that among the 150 members of treatment group 1, 30 percent responded to the email and among the 223 members of treatment group 3, 33 percent responded to the email. The difference in percentages between the two groups is small and not statistically distinguishable from zero.

Because we are interested in estimating the fraction of caregivers that were aware of the training, in addition to those who responded to the \$10 incentive, it is reasonable to assume that any individual beginning or completing the online training module was also aware of its availability. There were 20 individuals who were offered the \$10 incentive and did not respond to it, but either began or completed the training. These individuals comprise 5 percent of those offered \$10 to respond and are a small group relative to those who did respond to the \$10 incentive (20 vs. 119). As seen in Figure 4b, combining these types of responses leads us to estimate that 37 percent of individuals were aware of the training opportunity. This likely represents a lower bound since some individuals might have read the email, not responded to it, and not completed the training.

We can incorporate this information into our demand estimates in at least two ways: 1)



(a) Response Rates



(b) Fraction Participants Aware of Training Opportunity

Figure 4: Response Rates to \$10 Offer and Awareness of Training Opportunity

directly scale up participation by the fraction that was aware (‘Scaled’) and 2) estimate the fraction who participated among those who were aware of the opportunity (‘Aware’). Using these approaches, we estimate that participation in the training program among informed caregivers was approximately 35 percent to 51 percent when the price is zero and approximately 59 percent to 65 percent when participants are paid \$50 to complete the training. The corresponding demand curves, as well as the demand curve based on our initial estimates (‘Initial’), are displayed in Figure 3.

Because our estimates are coming from a small sample in one state, it is useful to compare them to the results in Hidalgo et al. (2014). In that paper, the authors studied an experiment in the Netherlands in which vouchers for training programs were randomized to a set of low-skilled individuals. On average, the €1,000 vouchers were spent on training programs that took approximately 48 hours to complete; this corresponds to a per-hour price reduction of €21.<sup>16</sup> After converting to 2021 real U.S. dollars, the authors’ results indicate that lowering the price of training by approximately \$30 per hour increased training completion by 44 percent, about 1.47 percent per dollar. We find broadly similar results – a \$50 price reduction increased training completion by 67 percent, about 1.34 percent per dollar – despite numerous differences in setting, context, and duration and content of training. So while our results are from a small group that is not likely to be perfectly representative of all caregivers in the United States, the comparison to estimates from such a different context provides some hope that our estimates have a reasonable amount of external validity.

With that in mind, there are a few takeaways from these demand curves. When the caretakers were aware of the opportunity, a large fraction were willing to complete the training without being compensated—somewhere between one-third and one-half of caregivers completed the training when it was free. Despite that, there appears to be a long tail of individuals that would have to be paid very large amounts of money to complete the training. Using our demand curves where everyone is aware of the training opportunity, our estimates imply that 100 percent participation requires that caregivers are paid between \$139 (‘Scaled’) and \$175 (‘Aware’). Of course, this is highly speculative since it is far outside of the range of our data, but it suggests that the implied costs of completing training vary tremendously across caregivers.

## 7 Welfare

A number of federal and state policies promote expansions in caregiver training. Massachusetts already offers the Personal and Home Care Aide State Training program, or

---

<sup>16</sup>This assumes that trainees spent the full voucher and paid nothing out-of-pocket for the training. Spending less than the full voucher would lead this to be an over-estimate of the price reduction. Spending more than the voucher and, thus, supplementing the cost of training out-of-pocket would lead this to be an under-estimate of the price reduction.

PHCAST, which features 11 free online training modules and a “certificate of completion” that is earned upon scoring an 80% or better on each module-specific quiz. More similar in scope to the training used in our experiment, states commonly use one-time online training modules to convey disease-specific (e.g., COVID-19) information to home caregivers. The welfare implications of such programs depend on numerous factors, including both the caregivers’ and patients’ tastes for training, the (potentially subsidized) price of the training, and the efficacy of the training program (i.e., whether the training improves quality of life for the patient and/or lowers healthcare costs). In this section, we use our experimental findings to speculate on the welfare implications of online training programs.

We study a policy that provides access to and subsidizes an online training program through the lens of the Marginal Value of Public Funds (MVPF). The MVPF measures the ratio of the marginal benefit of a policy to the net cost of the policy to the government.<sup>17</sup> In our setting, there are two primary groups of individuals who might benefit from the policy, the caregivers and those receiving the care. The net cost to the government includes both the mechanical costs of the policy as well as any fiscal externalities that might arise.

We parameterize the MVPF as

$$MVPF = \frac{\overbrace{\left[ \int_0^\theta D^{-1}(q; \alpha_4) dq + \alpha_2 \theta \right]}^{WTP_C} \theta + \overbrace{\alpha_1 \alpha_4 \theta}^{WTP_P}}{\underbrace{(\alpha_2 + \alpha_3) \theta}_{MC} - \underbrace{\alpha_4 \theta}_{FE}}. \quad (3)$$

Starting with the denominator, the mechanical cost to the government of training a single caregiver is  $\alpha_2 + \alpha_3$ , where  $\alpha_2$  is the financial incentive offered to caregivers to complete training and  $\alpha_3$  represents the marginal non-incentive costs to the government of training a caregiver.<sup>18</sup> The parameter  $\alpha_4$  measures the fiscal externality of training a single caregiver. There are many pathways through which training might lead to fiscal externalities: reducing Medicaid expenditures on care recipients, reducing caregivers’ reliance on other social programs, pulling workers from other jobs that are more/less subsidized than caregiving, etc. For simplicity, we will assume that the dominant fiscal externality to training is a potential reduction in the present discounted value of Medicaid spending that results from an efficacious training program. Both the mechanical cost and fiscal externality are then scaled by

<sup>17</sup>Because Finkelstein and Hendren (2020) describe the MVPF and its applications for welfare analysis clearly, we will only briefly outline it here. The tool has been around for much longer (e.g., Mayshar (1990)), but has garnered recent attention due to works like Hendren and Sprung-Keyser (2020).

<sup>18</sup>The latter cost will vary according to the government’s approach to creating, disseminating, and monitoring training applications. Throughout this section, we will assume the government utilizes a third-party for these purposes, such as the organization RCTCLEARN.NET that supplied training and monitored completion for our experiment. As RCTCLEARN.NET regularly contracts with home healthcare trade organizations and agencies, contracting with a state Medicaid office on a per-user basis at a relatively low cost would be well within their business model.

the share of caregivers completing training,  $\theta$ . Again for simplicity, by scaling all costs and benefits by  $\theta$ , we assume that each caregiver has just one patient and each patient just one caregiver.

In the numerator,  $WTP_C$  reflects a caregiver’s willingness to pay for the training. In many instances, it is reasonable to estimate the beneficiaries’ willingness to pay for training with an envelope theorem argument: subsidizing training by a small amount only impacts consumer welfare through direct effects, not through behavioral changes that result from the subsidy. However, in our case, the subsidy is discrete—we are providing access to a training program that was otherwise unavailable to the caregivers. Because of this, we can not appeal directly to the envelope theorem and instead need caregivers’ demand curve for training to estimate their willingness to pay for the policy.<sup>19</sup>

It is important to note that caregivers’ perceptions of the training play a key role in their demand. Prior to beginning the program, caregivers have beliefs about how training will affect the quality of the care they provide, how the training might affect their future labor market opportunities, how much they might enjoy or dislike the training (its consumption value), how it might affect their relationships with patients, and so on. These beliefs affect caregivers’ choices and thereby estimates of their willingness to pay for the training. To the extent that caregivers systematically overestimate positive impacts of training, the demand curve will overestimate their (ex post) willingness to pay for training; to the extent they underestimate positive impacts, the opposite is true.

We have some evidence that trainees are generally satisfied with the SBAR training module and find it useful. Upon passing the quiz at the conclusion of the module, trainees must complete a short survey prior to receiving a certificate of completion.<sup>20</sup> Ninety-nine percent of surveyed trainees report being “satisfied with the learning experience.” Fifty-two percent report that they “intend to use the knowledge gained from this course in [their] practice.” Trainees were also asked open-ended questions regarding what they liked about the training, what they didn’t like, and an example of how they might use the SBAR technique. We list representative responses in Appendix F. Generally, the responses indicate that the information presented was easy to understand and could be useful. This suggests that their expectations—which form the basis for their willingness to pay—might not have been overly optimistic. However, because these responses occurred right after the training, it does not rule out the possibility that the information learned is quickly forgotten or never used. In that case, the portion of ex-ante willingness to pay due to expected improvements in care

---

<sup>19</sup>There is an important distinction here between individuals who are currently caregivers and those who might become caregivers to gain access to the training if the policy were enacted. Our estimate of the MVPF focuses on current caregivers under the assumption that the latter group is likely to be small.

<sup>20</sup>This survey is administered by RCTCLEARN.NET and is given to all trainees, not just those from HRA. We were not able to customize this survey and were not allowed to match survey responses to individual trainees.

might still overstate the realized benefits from the specific training.

We let  $D(p; \alpha_4)$  represent caregiver demand for training. It tells us the fraction of caregivers who complete the training at each price,  $p$ . The  $WTP_C$  can then be measured as the sum of two pieces: (i) the integral of  $D^{-1}(q; \alpha_4)$  between zero and  $\theta$  and (ii) the financial incentive,  $\alpha_2$ , which is only received by the share of individuals that train.<sup>21</sup> Past work such as Barczyk and Kredler (2018) and Mommaerts (2020) has assumed that child caregivers internalize the utility of their parents, the latter of which is likely a function of parental health. To recognize this, we explicitly note that demand for the training is a function of  $\alpha_4$ , which can be thought of as a proxy for the impact of training on patient health.

We have written  $WTP_P$  as  $\alpha_1 \alpha_4 \theta$ , which assumes that the patient’s willingness to pay for their caregiver to train is a function of the efficacy of the training. For each dollar of health benefits  $\alpha_4$ , the patient is willing to pay  $\alpha_1$ . Like the denominator, the numerator needs to be scaled by the share of caregivers that train,  $\theta$ .

Equation (3) represents a welfare metric that can generalize to virtually any caregiver training incentive program. Moving forward, we use our experimental evidence to inform several elements in Equation (3). As such, the resulting MVPF is specific to dollars spent on a program that mimics our experiment; in particular, a program that notifies caregivers, via email, that a one-hour online training module has been made available at a particular (subsidized) price. Moreover, the calculations that follow assume that the ex-ante perceptions of the experimental training are representative of the perceptions caretakers would have regarding a hypothetical, Medicaid-supplied training. Regarding *this* MVPF, our experiment provides the key information needed to estimate  $D^{-1}(q; \alpha_4)$ .<sup>22</sup>

In addition to the information from the experiment, we need to make assumptions about other parameters. First, assume that the Texas Medicaid office offers training at a price of zero with no financial incentive (i.e.,  $\alpha_2 = 0$ ). It is likely that  $\alpha_3$  is quite small. According to the website of RCTCLEARN.NET, “the total price (of their services) can be less than a dollar per contact hour.” Under these conditions,  $MVPF = (4.81 + \alpha_1 \alpha_4)/(1 - \alpha_4)$ . Since that object is increasing in  $\alpha_4$ , a lower bound on the MVPF in this setting is 4.81; the bootstrapped 95-percent confidence interval centered on that lower bound is [1.99, 12.66]. The obvious next question is then, *is this large?* Because the MVPF does not account for the welfare costs of the government raising the funds to pay for a program, it is best suited for comparing welfare impacts across programs. In particular, if policy A has an MVPF of

---

<sup>21</sup>Note that this formulation is equivalent to  $\int_{p_s}^{\bar{p}} D(p; \alpha_4) dp$  where  $p_s$  is the subsidized price of training for the caregiver (which could be negative),  $D(p_s; \alpha_4) = \theta$ , and  $D(\bar{p}; \alpha_4) = 0$ .

<sup>22</sup>We use our ‘Initial’ demand curve in the analysis that follows, the functional form of which is presented in Appendix H. As stated above, this demand curve assumes linearity, similar to the welfare analysis in Einav et al. (2010). Our conclusions are robust to using alternative, non-linear demand curves (see Appendix G). The analysis could also be performed using the ‘Scaled’ or ‘Aware’ demand curves, which assume that all caregivers are aware of the training; however, such analysis would need to account for advertising costs when calculating ‘mechanical costs’ in Equation (3).

2 while policy B has an MVPF of 1 and both policies affect the same set of individuals, then we could raise welfare by taking one dollar of net spending away from policy B and transferring it to policy A.

To know whether the MVPF for the training program is large, we compare it to the MVPF for a hypothetical policy in which Texas’s Medicaid program raises caregiver wages. This alternative policy is motivated by regular calls for higher wages for home care workers, most recently in President Biden’s Build Back Better agenda. Based on our experimental results, offering the training at a price of zero would cost the Texas government approximately \$133,000. The cost-equivalent increase in Medicaid wages for caregivers is very small, no more than \$0.02 per hour (see Appendix I for the details of these calculations). We estimate the MVPF for a marginal increase in Medicaid’s caregiver-wages to be no larger than 1 (see Appendix J for details). Relative to this policy the MVPF of simply offering training at a price of zero appears to be large.<sup>23</sup>

Second, assume again that the Texas Medicaid office offers training at a price of zero with no financial incentive (i.e.,  $\alpha_2 = 0$ ). If the numerator of the MVPF is positive and the denominator is negative, which would occur if  $\alpha_4 > \alpha_3$ , then the policy “pays for itself” as the impact of the policy has a net positive effect on the government’s balance sheet. Because the mechanical cost to Medicaid for one completer,  $\alpha_3$ , is almost certainly very small, the policy of simply making training available could pay for itself with even a weakly efficacious training module. To illustrate, note that the average cost to Medicaid of an emergency room visit is \$420 (Moore and Liang, 2020). If  $\alpha_3 = 1$ , then training needs to prevent just one ER visit for every 420 people that complete training in order for the policy to pay for itself. In a sample of healthy adults, ER visits are rare, making this a tall order. However, the disabled patients in our sample average roughly one ER visit annually according to HRA records and many of these ER visits may be avoidable (Konetzka et al., 2012); thus, preventing one per 420 trained caregivers is conceivable.

## 8 Discussion and Limitations

In this article, we report results from an RCT we implemented in which personal care aides were provided access to an online training module and a random subset was offered \$50 to complete the training. Thirteen percent of the control group completed training, while the incentive increased completion in the treatment group by nine percentage points. When scaled to reflect the proportion of caregivers that were likely aware of the training opportunity, the incentive increased completion from 35 to 59 percent. These results suggest

---

<sup>23</sup>For mechanical costs ( $\alpha_3$ ) less than \$4, the MVPF remains above 1, while the cost-equivalent Medicaid wage increase remains almost trivially small (much less than \$0.10 per hour), leaving the wage increase MVPF unchanged.



that for training programs like the one used in our experiment, a policy that makes training freely available, and potentially offers a financial incentive for completion, can result in a majority of caregivers taking up the training without great expense. However, if the policy aims to have 100 percent participation, a large financial incentive is required. We use the MVPF framework discussed in Finkelstein and Hendren (2020) to speculate on the welfare impacts of a policy that mimics our experiment: access to a short online training program. That many caregivers are willing to complete training when it is free makes the welfare implications of a government simply offering training opportunities at a price of zero favorable. Our analysis suggests that for our setting, subsidizing training offers greater welfare gains than a cost-equivalent \$0.02 increase to the hourly wage that Medicaid pays caregivers. Moreover, the low marginal cost (to governments) of supplying online training programs makes the programs likely to pay for themselves. As Medicaid spends so much on home-based care, if the training program improves care-recipients' health even a small amount, the resulting reduction in healthcare costs is likely larger in magnitude than the costs of providing the training.

There are several limitations of our findings that are worth noting. First, potential trainees in our experiment—and in skills training programs more generally—do not perfectly know the true distribution of returns to any particular training. This can lead the demand curves to overstate (understate) the ex-post willingness to pay if trainees' beliefs are systematically too positive (too negative). As an extreme case, consider a training program that is completely useless or even harmful. If potential trainees do not realize this until after they have gone through the training, then the ex-ante demand curve (prior to training) will overstate willingness to pay relative to an ex-post full information benchmark. While this could be a problem in some training programs, it does not seem to be a huge concern in our experiment because of trainees' perceived usefulness of the training materials. Second, our MVPF analyses assume that the fixed costs of creating the training materials have already been paid and that the government can purchase training services at a low marginal cost. Currently, this assumption is reasonable, but if new materials need to be created and paid for by the government, the fixed costs of development would need to be incorporated into welfare calculations. Third, extrapolating our results to longer or recurring training sessions may be problematic. The training module that we studied took approximately one hour to complete and was not part of a recurring training regime (as many continuing education requirements are). Fourth, it is not obvious that the functional form assumptions we impose on the demand for training and which play a role in our MVPF calculations are correct. For example, we assume that demand is continuous through the price of zero while there is some evidence from the developing world that demand might jump at zero (e.g. Kremer and Miguel, 2007). Finally, the experiment was conducted in February and March of 2021, which was during the COVID-19 pandemic; in particular, at the tail end of the winter wave, when

many Americans were beginning to receive vaccinations. We acknowledge that pandemic related factors (e.g., lower overall mobility, fewer hours worked, a potentially greater need for cash, etc.) could impact the generalizability of our results. That said, at this point in time, all caretakers had access to the vaccine for several months and mobility trackers, such as Safegraph, show that mobility had returned to prepandemic levels in the Dallas-Fort Worth Metroplex.<sup>24</sup>

While our estimates have a number of limitations, we believe they represent a useful step forward. More broadly, considering the limitations of our estimates throws into relief the need for more work in the area that could inform the active policy debate around training home care workers.

## References

- Acemoglu, D. and J.-S. Pischke (1998). Why do firms train? theory and evidence. *The Quarterly journal of economics* 113(1), 79–119.
- Acemoglu, D. and J.-S. Pischke (1999). The structure of wages and investment in general training. *Journal of political economy* 107(3), 539–572.
- Banerjee, A. V., S. Chassang, S. Montero, and E. Snowberg (2020, April). A theory of experimenters: Robustness, randomization, and balance. *American Economic Review* 110(4), 1206–1230.
- Barczyk, D. and M. Kredler (2018). Evaluating long-term-care policy options, taking the family seriously. *The Review of Economic Studies* 85(2), 766–809.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of political economy* 70(5, Part 2), 9–49.
- Blom, B. (2019). Federal mandatory spending for means-tested programs, 2009 to 2029.
- Brown, J. R. and A. Finkelstein (2011). Insuring long-term care in the united states. *Journal of Economic Perspectives* 25(4), 119–42.
- Bruhn, M. and D. McKenzie (2009, October). In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics* 1(4), 200–232.
- Burgdorf, J., D. L. Roth, C. Riffin, and J. L. Wolff (2019). Factors associated with receipt of training among caregivers of older adults. *JAMA Internal Medicine* 179(6), 833–835.
- Cebul, R. D., J. B. Rebitzer, L. J. Taylor, and M. E. Votruba (2008). Organizational fragmentation and care quality in the us healthcare system. *Journal of Economic Perspectives* 22(4), 93–113.
- Chang, C. and Y. Wang (1996). Human capital investment under asymmetric information: The pigovian conjecture revisited. *Journal of labor Economics* 14(3), 505–519.

---

<sup>24</sup><https://www.safegraph.com/data-examples/covid19-commerce-patterns>

- Cooper, C., B. Cenko, B. Dow, and P. Rapaport (2017). A systematic review evaluating the impact of paid home carer training, supervision, and other interventions on the health and well-being of older home care clients. *International psychogeriatrics* 29(4), 595–604.
- Davis, D. A., M. A. Thomson, A. D. Oxman, and R. B. Haynes (1995). Changing physician performance: a systematic review of the effect of continuing medical education strategies. *Jama* 274(9), 700–705.
- Eiken, S., K. Sredl, B. Burwell, and A. Amos (2018). Medicaid expenditures for long-term services and supports in fy 2016. Technical report, IBM Watson Health.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics* 125(3), 877–921.
- Einav, L., A. Finkelstein, and N. Mahoney (2018). Provider incentives and healthcare costs: Evidence from long-term care hospitals. *Econometrica* 86(6), 2161–2219.
- Eiserer, T. (2019, May). Special needs families: 10 years too long to wait for state help. <https://www.wfaa.com/article/news/local/investigates/special-needs-families-10-years-too-long-to-wait-for-state-help/287-7b11d72e-39bf-497c-bd8b-79ca258d3666>.
- Finkelstein, A. and N. Hendren (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives* 34(4), 146–67.
- Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics* 135(3), 1209–1318.
- Hidalgo, D., H. Oosterbeek, and D. Webbink (2014). The impact of training vouchers on low-skilled workers. *Labour Economics* 31, 117–128.
- H.R.5376. 117th Congress (2021-2022): Build Back Better Act. <https://www.congress.gov/bill/117th-congress/house-bill/5376/text>, Accessed: 10/7/2021.
- Iritani, K. M. (2017, May). Medicaid personal care services, more harmonized program requirements and better data are needed. Technical report, United States Government Accountability Office. <https://www.gao.gov/assets/gao-17-598t.pdf>.
- Khatony, A., N. D. Nayery, F. Ahmadi, H. Haghani, and K. Vehvilainen-Julkunen (2009). The effectiveness of web-based and face-to-face continuing education methods on nurses’ knowledge about aids: a comparative study. *BMC Medical Education* 9(1), 1–7.
- King, A. I., M. Parsons, E. Robinson, and D. Jørgensen (2012). Assessing the impact of a restorative home care service in new zealand: a cluster randomised controlled trial. *Health & Social Care in the Community* 20(4), 365–374.
- Konetzka, R. T., S. L. Karon, and D. Potter (2012). Users of medicaid home and community-based services are especially vulnerable to costly avoidable hospital admissions. *Health Affairs* 31(6), 1167–1175.
- Kremer, M. and E. Miguel (2007). The illusion of sustainability. *The Quarterly journal of economics* 122(3), 1007–1065.

- LeBlanc, A. J., M. Tonner, and C. Harrington (2001). State medicaid programs offering personal care services. *Health Care Financing Review* 22(4), 155–173.
- Leuven, E. and H. Oosterbeek (1997). *Demand and supply of work-related training: Evidence from four countries*, Volume 18. Citeseer.
- Marinopoulos, S. S., T. Dorman, N. Ratanawongsa, L. M. Wilson, B. H. Ashar, J. L. Magaziner, R. G. Miller, P. A. Thomas, G. P. Prokopowicz, R. Qayyum, et al. (2007). Effectiveness of continuing medical education. *Evidence report/technology assessment* (149), 1–69.
- Mayshar, J. (1990). On measures of excess burden and their application. *Journal of Public Economics* 43(3), 263–289.
- Mommaerts, C. (2020). Long-term care insurance and the family. Technical report, University of Wisconsin Working Paper.
- Mommaerts, C. and Y. Truskinovsky (2020). The cyclicalities of informal care. *Journal of Health Economics* 71, 1–26.
- Moore, B. J. and L. Liang (2020). Costs of emergency department visits in the united states, 2017. In *Healthcare Cost and Utilization Project (HCUP) Statistical Briefs [Internet]*. Agency for Healthcare Research and Quality (US).
- Musumeci, M., P. Chidambaram, and M. O’Malley Watts (2019, April). Key questions about medicaid home and community-based services waiver waiting lists. Technical report, Kaiser Family Foundation.
- O’Malley Watts, M., M. Musumeci, and P. Chidambaram (2020, February). Medicaid home and community-based services enrollment and spending. Technical report, Kaiser Family Foundation.
- Osterman, P. (2017). *Who will care for us?: Long-term care and the long-term workforce*. Russell Sage Foundation.
- PHI (2016). U.s. home care workers: Key facts. <https://phinational.org/wp-content/uploads/legacy/phi-home-care-workers-key-facts.pdf>. Accessed 10/08/2021.
- PHI (2021). Personal care aide training requirements. <http://phinational.org/advocacy/personal-care-aide-training-requirements/>. Accessed 04/19/2021.
- Ruggles, S., S. Flood, S. Foster, R. Goeken, J. Pacas, M. Schouweiler, and M. Sobek (2021). Ipums usa: Version 11.0 [dataset]. <https://doi.org/10.18128/D010.V11.0>.
- Skira, M. M. (2015). Dynamic wage and employment effects of elder parent care. *International Economic Review* 56(1), 63–93.
- Spetz, J., R. I. Stone, S. A. Chapman, and N. Bryant (2019). Home and community-based workforce for patients with serious illness requires support to meet growing needs. *Health Affairs* 38(6), 902–909.

- Spetz, J., L. Trupin, T. Bates, and J. M. Coffman (2015). Future demand for long-term care workers will be influenced by demographic and utilization changes. *Health Affairs* 34(6), 936–945.
- Stevens, M. (1994). A theoretical model of on-the-job training with imperfect competition. *Oxford economic papers*, 537–562.
- Van Houtven, C. H., N. B. Coe, and M. M. Skira (2013). The effect of informal care on work and wages. *Journal of Health Economics* 32(1), 240–252.
- Van Houtven, C. H., V. A. Smith, J. H. Lindquist, J. G. Chapman, C. Hendrix, S. N. Hastings, E. Z. Oddone, H. A. King, M. Shepherd-Banigan, and M. Weinberger (2019). Family caregiver skills training to improve experiences of care: a randomized clinical trial. *Journal of general internal medicine* 34(10), 2114–2122.

# Appendices

## A Comparison of Sample to Formal and Informal Caregivers

In this section, we compare our sample’s demographics to those of other PCAs and informal caregivers. This helps us assess whether our sample population appears to be representative of two groups of interest: individuals whose profession is home-based caregiving and informal caregivers. For the former, we use data from the American Community Survey; for the latter, we use reported summary statistics from Mommaerts and Truskinovsky (2020).

There are a number of complications that make estimating the demographic characteristics of PCAs far less straightforward than it might seem. First, there are two groups of individuals who fall into the category of formally employed, paid, home-based caregivers and they can be difficult to distinguish from each other in the data: home health aides and PCAs.<sup>25</sup> Second, some (unknown) fraction of PCAs are self-employed and working in the gray market, not reporting all of their income to the IRS. With these difficulties acknowledged, we use data from the IPUMS USA (Ruggles et al., 2021) versions of the 2016 and 2017 ACS in conjunction with the procedures outlined in Osterman (2017) to identify PCAs.<sup>26</sup> For the demographics of informal caregivers, we pull statistics from Table 1 in Mommaerts and Truskinovsky (2020); these statistics are based on individuals who are not formally employed as caregivers as reported in the American Time Use Survey.

Table A1 reports the means of available demographic variables for our sample (column (1)), all PCAs in the ACS (column (2)), and informal caregivers (column (3)). The HRA sample is fairly similar to the overall population of formal PCAs in the ACS in terms of gender, race, age, regular hours worked per week, and wage rate. Because the means for informal caregivers are not available for all variables, it is less clear how similar or different the HRA caregivers are from this group. However, for the available variables, HRA’s caregivers are of a similar age to informal caregivers, but are considerably more likely to be female, to not be White, and to provide many more hours of care per week. Taken together, these summary statistics suggest that the HRA sample is more representative of formal PCAs nationally than it is of informal caregivers.

---

<sup>25</sup>Recall the distinction that home health aides provide medical care in addition to helping with ADL and IADL limitations, while PCAs only perform the latter. PCAs have different names in different areas including personal care assistants, personal care aides, home care aides, and others.

<sup>26</sup>We use the IPUMS occupation code 4610 and IPUMS industry codes 7580, 8170, 8180, 8370, 9290, and 9480. The years chosen are the most recent years for which the coding scheme outlined in Osterman (2017) is feasible.

Table A1: Demographics of Formal Caregivers, our PCAs, and Informal Caregivers

	HRA sample	ACS	Informal caregivers
	(1)	(2)	(3)
Female	0.80	0.84	0.54
Black	0.28	0.21	.
White	0.47	0.58	0.71
Hispanic	0.19	0.23	.
Age	43.71	45.93	45.90
Usual hours worked per week	30.32	32.61	6.07
Wage rate	12.16	12.15	.

Notes: Means provided for for our HRA sample (column (1)), all formal caregivers as described in the text in the ACS (column (2)), and informal caregivers (column (3)). Columns (1)-(2) are the authors' calculations; column (3) is based on Mommaerts and Truskinovsky (2020).

## B Text of Emails

In this section, we show the text of the emails included in the experiment. The emails came from “HRA Caregiver Services,” which is the same name that appears on each caregiver’s pay check and other communications between HRA and the caregivers. The subject line for all emails was, “Training Opportunity.” Elements included in double braces below, “{{ }},” were replaced with the relevant information for each individual caregiver. Deviations from the control group email are highlighted (but not highlighted in actual emails sent to caregivers). Detailed instructions on how to access the training module, referenced in the emails below, were attached to the emails.

*Control email:*

From: HRA Caregiver Services  
Subject line: Training Opportunity

Dear {{Name}},

We at HRA are offering you an opportunity to complete a **free caretaker training program**. The attached document provides simple steps to help you access and begin the training program using the token below. This opportunity will only be available until 11:59 PM on **Sunday, February 28**.

If you score 80% or better on the final quiz, you will be offered a printable certificate of completion. You may attempt the final quiz as many times as you please. We expect that the training will take a maximum of one hour.

**TOKEN: {{token}}**

We encourage you to contact us by email at [research@hratexas.org](mailto:research@hratexas.org) or by phone at **{{phone}}** with any questions you have about program access and completion.

Sincerely,  
{{Name}}  
{{email.com}}  
{{xxx-xxx-xxxx}}

*Treatment 1 email, \$50 incentive and \$10 response incentive:*

From: HRA Caregiver Services  
Subject line: Training Opportunity

Dear {{Name}},

We at HRA are offering you an opportunity to complete a **free caretaker training program**. The attached document provides simple steps to help you access and begin the training program using the token below. This opportunity will only be available until 11:59 PM on **Sunday, February 28**.

If you score 80% or better on the final quiz, you will be offered a printable certificate of completion **and we will mail you a \$50 Visa gift card by the end of March**. You may attempt the final quiz as many times as you please. We expect that the training will take a maximum of one hour.

**In addition to providing training information, we are trying to confirm that our emails are being received and read by caretakers. If you reply to this email with "got it", we will send you a \$10 Visa gift card.**

**TOKEN: {{token}}**

We encourage you to contact us by email at [research@hratexas.org](mailto:research@hratexas.org) or by phone at **{{phone}}** with any questions you have about program access and completion.

Sincerely,  
{{Name}}  
{{email.com}}  
{{xxx-xxx-xxxx}}



*Treatment 2 email, \$50 incentive:*

From: HRA Caregiver Services  
Subject line: Training Opportunity

Dear {{Name}},

We at HRA are offering you an opportunity to complete a **free caretaker training program**. The attached document provides simple steps to help you access and begin the training program using the token below. This opportunity will only be available until 11:59 PM on **Sunday, February 28**.

If you score 80% or better on the final quiz, you will be offered a printable certificate of completion and we will mail you a \$50 Visa gift card by the end of March. You may attempt the final quiz as many times as you please. We expect that the training will take a maximum of one hour.

**TOKEN: {{token}}**

We encourage you to contact us by email at [research@hratexas.org](mailto:research@hratexas.org) or by phone at **{{phone}}** with any questions you have about program access and completion.

Sincerely,  
{{Name}}  
{{email.com}}  
{{xxx-xxx-xxxx}}

*Treatment 3 email, \$10 response incentive:*

From: HRA Caregiver Services  
Subject line: Training Opportunity

Dear {{Name}},

We at HRA are offering you an opportunity to complete a **free caretaker training program**. The attached document provides simple steps to help you access and begin the training program using the token below. This opportunity will only be available until 11:59 PM on **Sunday, February 28**.

If you score 80% or better on the final quiz, you will be offered a printable certificate of completion. You may attempt the final quiz as many times as you please. We expect that the training will take a maximum of one hour.

In addition to providing training information, we are trying to confirm that our emails are being received and read by caretakers. If you reply to this email with "got it", we will send you a \$10 Visa giftcard.

**TOKEN: {{token}}**

We encourage you to contact us by email at [research@hratexas.org](mailto:research@hratexas.org) or by phone at **{{phone}}** with any questions you have about program access and completion.

Sincerely,  
{{Name}}  
{{email.com}}  
{{xxx-xxx-xxxx}}

## C An Earlier HRA Training Experiment

In Fall of 2020 we piloted an email experiment with HRA caregivers in which they were informed of a free training opportunity on RCTCLEARN.NET. The specific trainings were unrelated to the module used in our current experiment. We randomized both the messaging strategy used in the email (some caregivers received a standard email, some received an email designed to prime altruistic feelings for one’s caregiver, and some were told that their caregiver would be informed if the caregiver were to complete the training) and the financial incentive offered for completion (caregivers were paid \$0, \$80, or \$130 to complete roughly three hours of training). More details can be seen in the AEA RCT Registry (AEARCTR-0006553). We found suggestive evidence that financial payments increase the probability of completion, but little evidence that the messaging strategies had an impact. The experiment described in this manuscript was designed to address lower than expected statistical power in the pilot. In particular, fewer than expected caregivers in the initial experiment were aware of the training opportunity (which we learned via followup surveys). The sample for the experiment described in this manuscript only includes the 747 caregivers that were not offered financial incentives in the first experiment. Those previously receiving a financial incentive were excluded (i) to ensure that all included caregivers shared similar feelings regarding the credibility of our offer and (ii) because we worried the relative size of the financial offer could differentially affect choices. Because randomization was balanced in the first experiment, our estimation sample is representative of the HRA population of caregivers at large.

Table C1 shows regressions that test whether the previous experiment has affected our results in this experiment. Column (1) provides our baseline regression results. Column (2) includes indicator variables for each of the three possible treatments the caregiver could have received in the first experiment: 1) the altruism prompt, 2) contacting the patient if the caregiver were to complete the training, and 3) being in the control group. Including past treatment status has almost no impact on our estimated impact of being offered \$50 in the current experiment. Column (3) presents results in which the \$50 monetary treatment was interacted with each of the possible treatments from the first experiment. In each case, the estimated impact of receiving the \$50 offer appears to substantially increase the probability of completing the training, though the individual point estimates are not always statistically significant. However, we can not reject the null hypothesis that all three coefficients are equal to each other ( $p = 0.772$ ) and we can not reject that any two of the given three coefficients are equal to each other (p-values ranging from 0.516 to 0.904).

Table C1: Impact of the \$50 Offer on Beginning Training and the Past Experiment

	Baseline	Past Treatment Controls	Interacted Past Treatments
	(1)	(2)	(3)
\$50 to complete	0.087 (0.027) [0.001]	0.088 (0.027) [0.001]	
\$50 to complete * Altruism			0.078 (0.046) [0.087]
\$50 to complete * Contact			0.070 (0.049) [0.154]
\$50 to complete * Control			0.113 (0.043) [0.009]
Control mean	0.129	0.129	0.129
R-squared	0.281	0.287	0.288
Observations	747	747	747

Notes: Dependent variable is an indicator for whether the caregiver completed the training module or not. Column (1) is our baseline specification that includes additional controls for the patient’s gender, the caregiver’s hourly wage and hours worked, and controls for whether those variables are missing. Column (2) adds indicator variables for each of the three possible groups the caregiver could have been randomized to in the first experiment. Column (3) includes the new controls from column (2) and interacts our \$50 treatment with each of those past experiment groups. Robust standard errors shown in parentheses; p-values presented in brackets.

## D AEA Registry Details

We registered this study in the AEA RCT Registry (AEARCTR-0007156) on February 10, 2021, prior to randomization and launch. Below is an exhaustive list of the ways in which we deviated from our pre-analysis plan.

- We listed “registered token” as an outcome; however, this information turned out to not be available from RCTCLEARN. Instead we considered “began training” as an outcome in columns (1) and (2) of Table 3, which is a close approximation. It is possible that some individuals registered their token, but never began the course; however, both metrics capture the notion that some caregivers start on a path to completion but at some point stop after learning that completing training is more taxing than initially thought.
- The registry specifies samples of 224, 150, 223, and 150 for control and treatment

groups 1, 2, and 3, respectively. A miscommunication led treatment status to be switched for two groups and led to samples of 224, 150, 150, and 223.

- The registry states that we will use the number of caregivers in one’s network and number of patients the caregiver is tending to as control variables; however, we have coded these variables as binary as most of the variation is at the extensive margin.
- We did not list several control variables used in our final analysis in the registry. None affect our results. The omissions include: (1) whether the individual previously completed online training (we stated that this variable would be used in rerandomization, but did not list it as a control); (2) pre-analysis survey source (we stated that we would control for *whether* someone completed the survey, but not the method by which they completed); (3) missing indicators; and (4) hours and wages (we did not know we could access these variables at the time of the registry).
- The pre-analysis plan does not mention subgroup analysis by whether the caregiver served any patients with other caregivers (Table 5, column (2)), though the variable was listed as a control. Despite the omission, this robustness test seemed important, as the caregiver network could have informed, for example, someone in the control group that others were receiving monetary awards for completing training, potentially influencing the training decision. Moreover, the pre-analysis plan does not mention testing whether the one week extension of the experiment due to the Texas power outage (Table 5, column (5)), as this event could not have been anticipated.

The pre-analysis plan mentions subgroup analysis by whether the caregiver is an HRA employee (Table 5, column (3)) and whether the patient is insured by Amerigroup (Table 5, column (4)). We also planned subgroup analysis by age, tenure with patient, relationship to patient prior to caregiving, whether currently providing care to patients, race/ethnicity, gender, whether completed pre-analysis survey, and months working for HRA. Table D2 below contains these analyses. These findings have been relegated to an appendix because (i) we did not have strong priors on whether, or in which direction, treatment effects would vary across these subgroups and (ii) we do not have sufficient power to reject the null for some economically meaningful differences.

Table D2: Estimating the Impact of the \$50 Offer by Subgroup

	Age	Months with patient	Not fam-ily/friend	Not providing care	Not white	Hispanic	Not female	Completed survey	Months with HRA
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
T	0.118 (0.037) [0.001]	0.115 (0.070) [0.101]	0.109 (0.043) [0.012]	0.116 (0.037) [0.002]	0.108 (0.047) [0.022]	0.082 (0.038) [0.031]	0.063 (0.040) [0.122]	0.126 (0.045) [0.005]	0.122 (0.040) [0.002]
T*Age	-0.059 (0.054) [0.278]								
T*Months with patient		-0.052 (0.073) [0.476]							
T*Not family/friend			-0.079 (0.070) [0.265]						
T*Not providing care				-0.077 (0.050) [0.125]					
T*Not white					-0.076 (0.067) [0.254]				
T*Hispanic						-0.052 (0.079) [0.510]			
T*Not female							0.163 (0.092) [0.077]		
T*Completed Survey								-0.051 (0.056) [0.361]	
T*Months with HRA									-0.041 (0.049) [0.394]
R-squared	0.280	0.285	0.299	0.283	0.281	0.279	0.295	0.282	0.230
Observations	729	500	509	747	544	544	477	747	689

Notes: Dependent variable is whether the caregiver completed training or not.  $T$  is an indicator for the \$50 treatment. All models include the rerandomization variables, patient gender, whether patient had only one caregiver, whether caregiver had only one patient, caregiver's hourly wage and hours worked, and controls for whether those variables are missing. When the variable being tested for heterogeneous effects is continuous, we create an indicator for whether the caregiver is above the median value and use that in the analysis. Robust standard errors shown in parentheses; p-values presented in brackets.

## E Additional Figures and Tables

Table E1: Estimating the Impact of the \$50 Offer on Beginning Training

	Rerandomization Controls	Additional Controls
	(1)	(2)
\$50 to complete	0.084 (0.028) [0.003]	0.087 (0.028) [0.002]
\$10 to respond	0.013 (0.028) [0.636]	0.013 (0.028) [0.632]
Control mean	0.170	0.170
R-squared	0.262	0.267
Observations	747	747

Notes: Dependent variable is an indicator for whether the caregiver began the training module or not. Column (1) only includes the variables used in the rerandomization as controls. Column (2) adds controls for the patient's gender, the caregiver's hourly wage and hours worked, and controls for whether those variables are missing. Robust standard errors shown in parentheses; p-values presented in brackets.

Table E2: Comparing Estimates from OLS and Marginal Effects from Probit

	Began training			Completed training				
	OLS (1)	Probit (2)	OLS (3)	Probit (4)	OLS (5)	Probit (6)	OLS (7)	Probit (8)
Treat 1: \$50 + \$10	0.094 (0.036) [0.009]	0.083 (0.036) [0.021]			0.098 (0.034) [0.005]	0.091 (0.034) [0.008]		
Treat 2: \$50	0.116 (0.041) [0.005]	0.114 (0.038) [0.003]			0.116 (0.038) [0.003]	0.117 (0.035) [0.001]		
Treat 3: control + \$10	0.035 (0.034) [0.307]	0.048 (0.035) [0.164]			0.040 (0.032) [0.208]	0.058 (0.033) [0.076]		
\$50 to complete			0.088 (0.028) [0.002]	0.075 (0.026) [0.003]			0.087 (0.027) [0.001]	0.075 (0.024) [0.002]
\$10 to respond			0.012 (0.027) [0.663]	0.015 (0.026) [0.577]			0.017 (0.026) [0.523]	0.021 (0.024) [0.391]
Observations	747	747	747	747	747	747	747	747

Notes: Dependent variable is an indicator for whether the caregiver began the training module or not (columns (1) - (4)) or completed training (columns (5) - (8)). Regression specifications are either OLS or Probit; when probit, the average of the marginal effects for the sample is reported. Robust standard errors shown in parentheses; p-values presented in brackets.

## F Survey Responses

Those passing the quiz at the end of the SBAR module were asked to complete a survey. Below, we list several open-ended questions on the survey and a representative selection of responses.

### **What did you like about this learning experience?**

Modal response was some form of “easy to understand.”

“I especially like the SBAR approach. I identified with the fact that one has to know their audience and how they receive information.”

“Learning is a life process. Any course to allow growth in a job positions better worker greatly appreciated.”

“Was a simple and direct interface. Just enough words to get the point across instead of endless reading. Really liked the layout.”

“It taught me how to communicate more effectively, and efficiently.”

“Valuable information in my field of work”

### **How could this experience be improved?**

Modal response was some form of “no comment/no change”

“Having more classes”

“Better pictures”

“More examples using the SBAR tool”

### **Please give an example of how you might use the SBAR tool**

“Because my client has a communication barrier I will teach him as well as use the SBAR Communication Tool”

“Good communication skills to use with everyone, but especially with therapists, parents, and healthcare providers of the patient.”

“I am a family caregiver for my daughter, and I so sometimes need to provide updates to the nurses or physical therapist on progress or current events. Having this tool will help ensure I am providing information in a common form that we can both use/understand.”

“I plan to state the facts and what led up to the incident. Give my opinion as to what I think might be the solution.”



## G The MVPF with Alternative Demand Curves

In this section, we estimate the demand curve for training using alternative functional forms. While the set of possible functional forms for demand curves is large, the set of feasible demand functions is limited because our dependent variable is binary and we have only two prices at which we observe individuals' choices. Given the discrete choice nature of our outcome variable, the natural alternative models are the probit and logit. We estimate each of these models and using the estimates, predict the fraction of individuals who would take up the training at various price levels.<sup>27</sup> The resulting demand curves are shown in Appendix Figure G1.




Figure G1: Demand for Training: Alternative Functional Forms

The probit and logit demands are very similar to each other over the entire range shown. The linear, probit, and logit demand curves are quite similar once the price falls below zero, but there is a noticeable difference for positive prices: The probit and logit curves rise much more rapidly than the linear curve as the fraction of caregivers who take up training falls below ten percent. This difference suggests that the caregivers' willingness to pay for training under the probit and logit demands will be larger than under the linear demand. As seen in

---

<sup>27</sup>More specifically, once we have the model's estimated parameters, for each price, we predict each caregiver's probability of taking up training and average across caregivers. We imposed that the fraction who take up is 12.9 percent at the price of zero to match the control group mean. This required a (small) vertical shift of each demand curve.

Appendix Table G3, at a price of zero, caregivers’ estimated willingness to pay for training is about 70 percent greater for probit and logit demand than it is for linear demand.

Table G3: Willingness to Pay with Alternative Models of Demand

	(1)
Linear	4.8
Probit	8.1
Logit	8.1

Notes: Column (1) shows caregivers’ willingness to pay for a training program with a price of zero. Each row presents the estimate for the stated model of demand.

## H MVPF Algebra

We found that 13 percent of control individuals (i.e., those paid nothing) completed training. Adding our estimated treatment effect of being eligible for the incentive, the incentive raised completion rates to 22 percent.<sup>28</sup> Assuming linearity, we can write the inverse demand curve as

$$D^{-1}(\theta) = 74 - 573 * \theta. \quad (4)$$

Solving for  $\theta$ , we get a function linking the share of caregivers that will complete training to the financial incentive offered to them

$$\theta = 0.13 + \frac{\alpha_2}{573}. \quad (5)$$

Plugging these two expressions into Equation (3) and evaluating the integral produces.

$$MVPF = \frac{4.81 + .13\alpha_2 + \frac{\alpha_2^2}{573*2} + \alpha_1\alpha_4}{\alpha_2 + \alpha_3 - \alpha_4}. \quad (6)$$

## I Estimating the Cost-Equivalent Medicaid Wage Increase

We begin by estimating the cost of a policy in which the state of Texas offers the training at a price of zero. From O’Malley Watts et al. (2020), we know that there are 516,600

<sup>28</sup>Throughout this section, we round our estimates and all numbers based on them; values and figures in the paper are based on machine precision.

individuals receiving HCBS in Texas in a given year. If each care recipient has two caregivers (on average), this translates to 1,033,200 caregivers. Our results suggest that 12.9% of those caregivers take up the training when  $\alpha_2 = 0$ ; this means 133,282 caregivers take up training. Because we have assumed that the per-user fee the government pays is \$1, the cost to the government of this policy is  $\approx \$133,282$ .

To our knowledge, there is no source which separates out Medicaid’s expenditures on caregivers from its expenditures on other home and community-based services (HCBS). Instead, we estimate the Texas Medicaid wage increase for caregivers with the same cost as the training policy (\$133,282). The key piece of information that we do not have is the number of hours of care provided by Medicaid’s caregivers. We proceed to first estimate an upper bound on the likely size of the Medicaid wage increase (which we use in our main text) and then provide a more realistic estimate based on additional assumptions.

In 2018, more than 99.8 percent of Texas’s total HCBS expenditures (\$7,397,800,000) were spent on programs which included caregivers (see Appendix Table 2 in O’Malley Watts et al., 2020). If all of these expenditures were attributed to caregiver wages and the average wage rate matched our sample, this would imply that there were  $\$7,388,300,000 / \$12.16 = 607,590,461$  hours of care provided. Given this number of hours, the cost of the training policy would imply a cost-equivalent increase in Medicaid’s wage rate for caregivers of  $\$133,282 / 607,590,461 = \$0.0002$  per hour. If instead only 1 percent of these expenditures were due to wages, then the implied hours of care falls to 6,075,905 and the implied increase in Medicaid wages is \$0.02 per hour. While it is extremely unlikely that only one percent of Medicaid HCBS expenditures are going to caregivers, the increase in Medicaid wages implied by this assumption is still extremely small in absolute terms. Because of this, we will use this bound in the paper when forming the MVPF for this wage increase policy.

We can form a potentially more realistic estimate of the Medicaid wage increase with additional assumptions. There are two broad authorities through which states are able to provide Medicaid HCBS: (i) state plan services and (ii) waivers. In Texas a little more than half of Medicaid HCBS spending is provided through state plan services (O’Malley Watts et al., 2020). Broadly, the types of services provided through the two authorities overlap to a large degree. For state plan services, approximately 86 percent of expenditures went towards “personal care services,” a program where the vast majority of expenditures are likely due to caregiver wages. If we assume that 86 percent of expenditures are due to caregiver wages more generally, then the cost-equivalent implied increase in Medicaid’s wage rate for caregivers is  $\$133,282 / 525,457,485 = 0.00025$ .

## J MVPF for Wage Increases

In this section, we calculate the MVPF for a marginal increase in the hourly wage rate paid by Medicaid to caregivers. Appealing to the envelope theorem, a caregiver’s willingness to pay for a marginal per hour wage increase would simply be the total number of hours she currently works,  $H$ , scaled by the tax rate.<sup>29</sup> This represents the increased income she would receive from the wage increase.

The denominator of the MVPF is somewhat more complicated because it must take into account all of the policy’s effects on the government’s budget constraint. The most direct effects on the government’s budget constraint are the mechanical costs of the program,  $H$ , less the tax revenue it receives from the caregiver’s additional  $H$  of income. Although any changes to hours worked do not enter the caregiver’s willingness to pay (because of the envelope theorem), they do potentially affect the government’s budget constraint. The two primary dimensions we consider are whether (i) the number of hours spent providing care changes and (ii) labor supply to other wage-earning (and taxed) jobs are affected. For example, labor supplied to caregiving might rise with the increased wage while labor supplied to other jobs might decrease. In that case, the taxes the government receives from additional labor supplied in caregiving are offset to some degree by lost taxes in the caregiver’s other job. Previous research suggests that the elasticity of hours worked outside of caregiving with respect to the caregiving wage is likely negative, but fairly small in magnitude (Skira, 2015).<sup>30</sup> For simplicity, we will assume that any increases in taxes due to additional caregiving hours are just offset by reductions in taxes due to reduced non-caregiving labor supply.

In addition to these intensive margin changes, the extensive margin changes are also important to consider. Raising Medicaid caregiver wages might induce some individuals to become caregivers. While these people do not affect the numerator of the MVPF (again, due to the envelope theorem), they might be switching from a job that is not paid for by the government. In this case, the government’s budget constraint should reflect the wages,  $W$ , it must now pay to this type of individual. Potentially counteracting this latter fiscal externality is the possibility of higher quality care. Individuals who become caregivers at

---

<sup>29</sup>For example, consider a model in which the consumer chooses consumption and leisure subject to a budget constraint:  $\max_{c,l} U(c,l) \quad s.t. \quad w(T-l)(1-\tau) + A = c$  where  $w$  is the hourly wage rate,  $T$  is total time,  $l$  is leisure,  $T-l = H$  is hours worked,  $\tau$  is the tax rate, and  $A$  is assets. If  $W$  is the solution to this problem and the marginal utility of income is  $\lambda$ , then the envelope theorem tells us that  $(dW/dw)/\lambda = H(1-\tau)$ .

<sup>30</sup>Skira (2015) does not directly estimate this cross-price wage elasticity, but it can be (roughly) inferred from the estimates contained therein. To be explicit, consider the papers “policy experiment 3” in which the individual is paid \$18,250 for intensively caring for another individual over a two year period. Because providing intensive care is defined to be 1,000 or more hours per year, this translates to an increase in the hourly wage of no more than \$9.13. That \$9.13 increase in the hourly wage led to an estimated 3% increase in the fraction of people not working at all, a 2% reduction in the fraction of people working part time, and a 12% reduction in the fraction of people working full time. Our policy would only raise the wage by \$1 per hour, reducing those estimated effects by roughly an order of magnitude.

a higher wage might provide higher quality care and so reduce their patients' health care expenditures (relative to the average quality of care previously provided by caregivers). We denote these reductions as  $Q$ .

Given these assumptions, we can estimate the MVPF for a marginal increase in Medicaid's wage for caregivers to be

$$MVPF^{wage} = \frac{H(1 - \tau)}{H(1 - \tau) + W - Q}.$$

When no individuals choose to become caregivers in response to the increased wage, the MVPF is one; when they do, the size of the MVPF depends upon whether the additional wages paid to new caregivers are larger or smaller than the reduced health expenditures from having a higher quality pool of caregivers. Because the cost-equivalent increase in Medicaid wages is so small (no more than \$0.02 per hour), changes on the extensive margin are likely to be negligible and that the MVPF for the wage increase is approximately equal to 1.