

Using Insurance Before You Lose It: Health Care Consumption at the End of Coverage*

Tal Gross[†]
Mailman School of Public Health
Columbia University

October 2010

Abstract

Transitions from one insurance program to another—or from insured status to uninsured status—are common. How these transitions affect individuals depends, in part, on whether consumers anticipate the loss of insurance. Potentially, consumers may “stock up” on health care before losing coverage. Such anticipatory consumption would demonstrate whether individuals are forward-looking in their demand for medical care. This paper studies one transition in insurance status, as teenagers move from their family’s coverage to uninsured status. It finds no evidence that individuals stock up on medical care before coverage ends, but rather a general decrease in health care consumption in the last month of coverage. The results imply that models of myopia may be relevant in the case of health care consumption.

Keywords: insurance coverage, health care consumption

JEL Classification: I11, I18, G22.

*I am grateful to David Autor, Michael Anderson, Joshua Angrist, Carlos Dobkin, Esther Dufo, Matthew Notowidigdo, Jonathan Gruber, and Nirupama Rao for helpful comments.

[†]e-mail: tg2370@columbia.edu.

1 Introduction

Insured Americans obtain health insurance through their employer, from a family member's employer, from public sources, or in the individual market. As a result, transitions in insurance coverage are common. Individuals may switch health insurance programs when they move from one employer to another, when they become unemployed, or when they become ineligible for their current health insurance plan. This makes insurance status not simply a cross-sectional issue but a longitudinal one; more individuals are uninsured for a brief period of time than are uninsured at any one point in time. Levy (2007), for example, estimates that roughly 30% of young adult males are currently uninsured, but that 55% are uninsured within a two-year interval. Indeed, transitions into uninsured status have become more common, despite large expansions of publicly-provided health insurance (Cutler and Gelber, 2009).

This paper analyzes one such transition in insurance status. It tests whether individuals anticipate losing health insurance and “stock up” on care in the weeks before coverage is lost. When faced with a loss of coverage, a forward-looking consumer recognizes that the price she faces for health care will soon increase. Before coverage ends, she may consume elective care and medication in order to prepare for that increase.

There exists anecdotal evidence that individuals do indeed stock up on care before coverage ends. Health care professionals at one four-year-college observed a 13% increase in visits to the pharmacy at the end of the academic year, and cite students graduating and becoming uninsured as the prime cause (Twiggs, 2009). Despite such anecdotal evidence, statistical evidence is limited. To my knowledge, Long et al. (1998) present the only statistical analysis of anticipatory consumption for individuals about to become uninsured. The authors examine the health care consumption of respondents to the Survey of Income and Program Participation who later lose

insurance coverage. The authors find no evidence that those respondents consume more health care in the months before they lose coverage than respondents who retain coverage. But the authors cannot isolate households who lose coverage exogenously. Moreover, given the limited information on health care consumption in the household survey, the authors cannot measure the path of consumption in the weeks until coverage is lost. They rely instead on self-reported doctor visits and hospitalizations.

Whether individuals anticipate the loss of insurance and stock up on medical care in advance is an important question for at least three reasons. First, forward-looking behavior affects how economists evaluate insurance transitions. Anderson et al. (2009) analyze the transition from insured status to uninsured status at age nineteen, as teenagers lose coverage under their family's plans. Such a transition can be anticipated, and thus to fully understand the effects of the transition, one needs to evaluate whether health care consumption changes *before* the transition to uninsured status. The same may apply to other research on health insurance in which cross-sectional patterns are studied but "anticipatory" consumption is not.

Second, such transitions in insurance status may lead to inefficient consumption of health care. Health is a stock variable with an optimal investment path (Grossman, 1972). Transitions from one insurance plan to another bring about changes in price that are neither based on the cost of production nor designed to prevent moral hazard. If investment in health exhibits diminishing returns in any one time period, then investments ought to be made smoothly over time. But the "churn" brought on by insurance transitions can induce a consumer to deviate from the efficient path of investment. For that reason, health insurance transitions present a potential efficiency loss, solely in their distortion of health care consumption.

Finally, behavior at the end of coverage can demonstrate whether individuals are forward-looking in their demand for medical care, and the degree to which individuals

are forward-looking is pivotal in research on the demand for health care. Consider, for instance, research on high-deductible health insurance plans. Such research is complicated by the issue of forward-looking behavior. Upon observing consumption of health care in a high-deductible plan, the analyst does not know which price consumers take as relevant. In theory, forward-looking agents will predict the probability that their annual spending will exceed the deductible, and choose their consumption based on the expected price. A woman who is six months pregnant at the start of the year will realize that expenses associated with childbirth will exceed the deductible. She will then consider health care at the beginning of the year to be free, even though, nominally, she is charged a co-payment. But in reality, consumers may act myopically, and focus exclusively on the contemporaneous price of medical care.

Keeler and Rolph (1988) grapple with this issue when estimating the price elasticity of demand for medical care from the results of the RAND Health Insurance experiment. They estimate a structural model and conclude that individuals react to deductibles myopically, responding only to the current price of medical care, rather than the expected, year-end price. They then assume myopic behavior in order to calculate a widely-cited elasticity of demand. This assumption of myopia is criticized by Kowalski (2009), who assumes perfect foresight. To my knowledge, though, no studies have conclusively demonstrated whether myopia is an appropriate assumption, or rather, whether individuals are forward-looking in their demand for care.

This paper focuses on the loss of coverage in order to test for forward-looking behavior. Several empirical challenges complicate any analysis of individuals losing coverage. Insurance transitions are not assigned at random. Individuals may experience health shocks that both end their employment and affect their demand for health care. Moreover, health insurance transitions are often voluntary. Economic theory suggests that agents will switch health care plans whenever faced with new

information on the distribution of their potential health care costs. To analyze health insurance transitions, one needs to isolate individuals who face a transition for reasons unrelated to their expected health care costs.

For that reason, this paper focuses on teenagers who lose coverage under their family’s plans. Insurance contracts cover dependents “eighteen and under,” and only cover older dependents who are full-time students. As a result, 5–10 percent of American teenagers lose health insurance coverage at age nineteen (Anderson et al., 2009). This paper uses that transition in order to learn how individual behavior changes as health insurance coverage comes to an end.

An advantage of focusing on nineteen-year-olds is that the population is likely to be in transition to uninsured status. A disadvantage is that teenagers are not representative consumers of health care. They are generally healthier than the general population, and teenagers who lose coverage at age nineteen are generally wealthy enough to have been privately insured, but not to have enrolled in college by age nineteen. Nevertheless, the econometric framework described below captures a consistent estimate of the change in behavior towards the end of insurance coverage. I discuss how the findings may relate to the general population.

The paper proceeds as follows. Section 2 introduces my empirical framework and discusses the sample restrictions required. Section 3 reviews the main results on health care consumption. Section 4 extends the basic results to consider alternative explanations. Section 5 concludes.

2 Empirical Approach

An examination of health care consumption towards the end of coverage requires detailed data on consumption for a sample of individuals who lose health insurance

exogenously. This section describes the data I use and the empirical framework designed to test such a hypothesis.

2.1 Data and Sample Restrictions

The 2003 Medstat MarketScan research database contains insurance claims for the universe of beneficiaries from a sample of organizations. The MarketScan database contains a record for each claim charged to the insurer for three types of expenditures: outpatient, inpatient, and drug expenditures. The date of each claim is recorded along with the monetary cost of the service. The MarketScan database is well-suited for this application. The 2003 database is large, containing the complete expenditure records of over six million individuals. It contains the precise date on which individuals lose coverage, and there is likely little measurement error, as the data are critical to each insurance company's billing procedures.

At the same time, the Medstat data present several drawbacks. The data contain limited demographic information aside from age, gender, and geographic location. More importantly, the data do not indicate why individuals lose coverage. There exists no variable in the data set that categorizes individuals by whether their loss of insurance is exogenous and anticipated. For that reason, I apply the following selection criteria in order to isolate a sample likely to be losing health insurance in an anticipated and exogenous manner.

1. The individual is listed as a dependent, rather than being the primary beneficiary or the spouse of the primary beneficiary.
2. The individual either loses coverage on the last day of a month (the treated sample) or does not lose coverage (the control group).¹

¹Of those dependents who lose coverage, over 99% are listed as losing coverage on the last day of a month. I eliminate the tiny proportion that lose coverage during the month. The analysis below

3. If the individual loses coverage, the primary beneficiary and spouse maintain coverage.
4. The individual begins the year at age 18.

These criteria narrow the sample to young adults likely to be losing their family's coverage and becoming uninsured. A drawback to the MarketScan database is that I cannot isolate a sample in which all of the individuals become uninsured. Surely, a portion of the sample loses coverage but gains insurance through another provider. Nevertheless, Collins et al. (2003) estimate that the share of teenagers who are uninsured triples from age eighteen to age nineteen. Relatively few eighteen-year-olds transition from their family's coverage to other coverage, and even fewer transition to coverage that is *more* generous than their parents' plan.

Table 1 offers evidence on how selection criteria 1 through 3 interact with age. The second and third columns display simple counts of observations for dependents that begin the year at the given age. The second column lists counts of dependents who either do not lose coverage or lose coverage with the primary beneficiary. These counts decrease with age; fewer young adults qualify as dependents as they age. The third column lists counts of observations who lose coverage during the year while the primary beneficiaries do not (the treated sample). The number of treated individuals jumps to 17,145 for individuals who begin the year at age eighteen (as compared to only 2,309 at age 17 and 8,007 at age 19). This spike at age 18 suggests that the selection criteria are indeed isolating teenagers affected by the variation described by Anderson et al. (2009). Below, I isolate this sample and track their consumption before coverage ends.²

is based on cell means for individuals ending coverage on the same day. Retaining individuals who lose coverage during the month would lead to hundreds of additional cell means based on only a handful of individuals for each cell.

²The number of treated individuals drastically increases at age 22. That pattern is likely driven

Table 2 presents descriptive statistics for those observations who satisfy selection restrictions 1 through 3 and also begin the year at age 18. The sample consists of 100,882 teenagers, 17,001 who lose coverage during the year. Kowalski (2009) estimates that 40% of individuals consume zero health care in a given year, and table 2 demonstrates that this mass point becomes even more pronounced for teenagers. Nearly 30% of the sample consumes no health care at any point during the year. For the sample that loses coverage, 55% consume no health care while still eligible. This pronounced mass point at zero presents an empirical challenge. Small changes in the rate of consumption are difficult to isolate when so few individuals consume care. In what follows, I implement several strategies as a result.

The last two columns of table 2 present average expenditures and consumption rates for the entire individual-day sample. The 17,001 individuals who lose coverage and 83,881 individuals who do not comprise a “micro” data set of over 33 million individual-day observations. Each individual in the control group is observed for 365 days, and each treated individual is observed for the days of the year until their coverage ends. Approximately one percent of the observations consume care on any given day, with an average expenditure of roughly forty dollars.

The dollars spent on care are a much less reliable outcome than the rate of consumption. The dollars spent are often not directly chosen by the consumer. For example, the RAND health insurance experiment found that more generous insurance plans lead to more emergency department visits, but that the dollars spent at the emergency department was not affected (Newhouse and Group, 1993). Once patients choose to consume a given treatment, the intensive margin is often not discretionary. Also, the dollars spent on care are extremely skewed. Only 366 observations in the treated group consume an inpatient visit, but the median expenditure is \$59,568 and

by rates of four-year college attendance. Upon exiting college, young adults often lose eligibility under their parents’ coverage, as they are no longer full-time students (Collins et al., 2003).

the mean is \$767,531. Since inpatient care is dramatically more expensive than other types of care, a very small share of individuals—less than one in ten thousand—are responsible for average expenditures of roughly \$50. For that reason, I focus on the share of individuals who consume care rather than the dollars spent.

2.2 Regression Framework

Given this sample of consumers, I estimate the anticipatory effects of losing insurance coverage. To do so, I compare individuals who will lose coverage in several weeks to individuals who lose coverage much later or do not lose coverage at all. Though the sample is selected, the comparison between such sub-groups captures the parameter of interest, so long as the date of loss of coverage is exogenous.

For each group of individuals who lose coverage in month e , I observe the percent who consume care on day t . Denote this value as $\bar{y}_{te} \in [0, 100]$. I run simple difference-in-difference regressions of the form:

$$\begin{aligned} \bar{y}_{te} = & \alpha_0 + \beta_1 \cdot I\{e = t + 1\} + \beta_2 \cdot I\{e = t + 2\} + \dots + \\ & \alpha_t + \alpha_e + \varepsilon_{te}. \end{aligned}$$

In this regression equation, the indicator functions, $I\{\cdot\}$, measure how long till a group loses coverage. The α_t fixed effects capture any seasonal movement in medical consumption. The α_e fixed effects capture any fixed differences between groups that lose coverage early in the year versus late in the year. The hypothesis that individuals stock up on health care before coverage ends implies that $\beta_j > 0$.

3 Main Estimates of the Change in Behavior Towards the End of Coverage

Table 3 presents estimates of the regression equation above. The first column presents results for all types of health care and the point estimates are plotted in figure 1. The results indicate a *drop* in consumption in the last month of coverage. Individuals are roughly 0.07 percentage points less likely to consume medical care in the three weeks before the week that coverage is lost. In the last month of coverage, patients gradually cut back on care. But more than four weeks before losing coverage, there is no significant change in the probability of consuming care. These point estimates are not statistically significant at conventional levels. But the 95% confidence intervals rule out increases in the probability of consumption above 0.05 percentage points. The regressions provide little evidence of stocking up.

The remaining columns in table 3 present these results by type of care. It is difficult to judge the magnitude of these estimates, since so few individuals consume care on any given day. For instance, only 0.01 percent of individuals have an inpatient stay on any given day. Nevertheless, columns 2 and 3 suggest that in the last month of coverage, the probability of consuming medication or an outpatient visit decrease by 0.05 percentage points, a roughly 10 percent change in the rates of those types of care.

Consider next health care consumption in dollars. Table 4 presents estimates of the regression equation above for mean expenditures by category of care. As argued above, expenditures are a far less informative outcome than consumption rates. It is thus unsurprising that the estimates in table 4 are no more precisely estimated than those in table 3. Nevertheless, the table generally demonstrates a similar pattern. Mean expenditures seem to decrease in the last month of coverage. Expenditures in the penultimate month of coverage are imprecisely estimated and do not suggest a

clear pattern.

4 Robustness Checks

This section offers several variations on the main specification above. A key drawback to the approach above is that so few individuals in the sample consume any care. Potentially, it is individuals with chronic conditions and a greater need for care who would consume care immediately before losing coverage. One approach to select that sample is to select individuals who consumed care at some point in the past. That selected sample would exhibit a higher rate of consumption, and potentially, lead to more precise estimates of the consumption path before coverage is lost.

To pursue that approach, I select only individuals who consume a positive amount of health care in January of 2003. I assume that health care consumption in January is not a reaction to the (then distant) loss of coverage. To make this assumption more palatable, I restrict the sample to individuals who lose coverage in May through November or who did not lose coverage. Thus, I assume that individuals may change their consumption in the 3–5 months before they lose coverage, but they do not consume differently over six months prior to losing coverage.

When I impose those restrictions on the sample, the average rate of consumption increases from 1.2 percent to 3.2 percent. This selected sample contains individuals much more likely to consume health care. Thus if the point estimates in table 3 are driven by the large share of zeroes, then this selected sample ought to be robust to that concern.

The second column of table 5 presents results for this selected sample. Nearly all point estimates for the last month of consumption are still negative. For instance, in the week before coverage is lost, individuals are half a percentage point less likely to consume medical care (and this point estimate is statistically significant). The

magnitude of that decrease in consumption seems to be driven equally between medication and outpatient services. Restricting the sample by previous consumption, in this way, does not change the basic conclusion of table 3.

The remaining columns of table 5 explore other, alternative explanations. Most importantly, I test whether the results above are driven by the socio-economic status of the dependents in the sample. A general concern with the results reported in section 3 is that I cannot isolate a sample of dependents that become uninsured. I have argued above that few dependents transition onto *more* generous insurance at age eighteen. But the specific nature of the transitions are difficult to ascertain. Potentially, the regressions above isolate a decrease in consumption during the last month of coverage solely because some individuals are indeed transitioning onto more generous coverage, and thus postponing consumption.

In order to evaluate this alternative explanation, I use the limited demographic information contained in the Medstat sample. I stratify dependents by a proxy for their probability of remaining insured. The Medstat sample lists the county of residence for the main beneficiary. I match each county to the share of the county's residents aged 25 and older who have attended at least some college, a statistic reported in the 2000 decennial census. I consider this statistic, "share college," an imprecise proxy for the probability that the given dependent will attend college.

The sample of dependents that lose coverage have a mean county share college of 48%. The control sample has a mean share college of 52%. This differential suggests that the proxy captures the significant socio-economic differences in between the samples. The variation at age nineteen in insurance coverage is primarily driven by the requirement that the dependent be a full-time student (Collins et al., 2003). Thus it is re-assuring that dependents who actually lose coverage have a higher probability of losing coverage, as predicted by the demographics of their home county.

Based on this proxy, I split the dependents in the sample into those with high and low probability of attending college.³ Columns 3 and 4 of table 5 present results of equation 1 estimated for these two samples. The results do not differ qualitatively between the two groups; both samples exhibit similar decreases in consumption in the last month of coverage, and no significant increase in the penultimate month of coverage. I tentatively conclude that prospects after coverage ends are not responsible for the consumption patterns above.

Finally, the last two columns of table 5 present results for males and females separately. Health care consumption often differs by gender (see, for instance, Billings et al. 2000), but these last two regressions suggest no significant difference by gender in consumption patterns at the end of coverage.

5 Discussion

Forward-looking agents ought to increase consumption as the end of coverage approaches. This paper, however, finds no evidence of such anticipatory consumption. It demonstrates a slight decrease in health care consumption in the last month of coverage, and no significant increase beforehand. The decrease in consumption may be driven by individuals awaiting a transition to more generous coverage. But stratifying the sample by a proxy for the probability that they maintain insured status does not significantly alter the results.

A more likely explanation for the pattern above is that in the last month of coverage, consumers are uncertain over when precisely they become uninsured. Consequently, consumption gradually decreases throughout the last month. Still, consumers could stock up on health care in the penultimate month of coverage, and this paper has found no evidence for that behavior.

³I categorize individuals as having a low probability of attending college if their primary beneficiary lives in a county in which less than 50% of its adult residents attended some college.

Clearly, more research is necessary in order to understand the degree to which consumers of health care are forward-looking. Future research may focus on older populations more likely to consume care, and potentially more likely to exhibit anticipatory consumption. But the evidence presented above suggests that myopia may indeed be a reasonable assumption.

References

- Anderson, M., C. Dobkin, and T. Gross (2009). The effect of health insurance coverage on utilization of health care services. Mimeo.
- Billings, J., N. Parikh, and T. Mijanovich (2000, November). Emergency department use: The New York story. Issue Brief, The Commonwealth Fund.
- Collins, S. R., C. Schoen, J. Kriss, M. Doty, and B. Mahato (2003). Rite of passage? Why young adults become uninsured and how new policies can help. Issue brief, The Commonwealth Fund.
- Cutler, D. and A. Gelber (2009). Changes in the incidence and duration of periods without insurance. *The New England Journal of Medicine* 360(17), 1740–1748.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *The Journal of Political Economy* 80(2), 223–255.
- Keeler, E. B. and J. E. Rolph (1988). The demand for episodes of treatment in the health insurance experiment. *Journal of Health Economics* 7, 337–367.
- Kowalski, A. (2009). Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care. Mimeo.
- Levy, H. (2007). Health insurance and the transition to adulthood. In S. Danziger and C. Rouse (Eds.), *The Price of Independence: The Economics of Early Adulthood*, pp. 84–106. Russell Sage Foundation.
- Long, S. H., M. S. Marquis, and J. Rodgers (1998). Do people shift their use of health services over time to take advantage of insurance? *Journal of Health Economics* 17(1), 105–115.
- Newhouse, J. P. and T. I. E. Group (1993). *Free for All? Lessons from the RAND Health Insurance Experiment*. Harvard University Press.
- Twiggs, E. (2009, April). Graduating seniors head to UHS for one last free check up. *The Michigan Daily*.

Table 1: Sample Restrictions by Age

Age	Do Not Lose Coverage, N	Loses Coverage, N	Percent Losing Coverage
15	109,426	861	0.8
16	110,275	968	0.9
17	113,019	2,309	2.0
18	93,297	17,145	15.5
19	80,823	8,007	9.0
20	74,326	7,170	8.8
21	62,459	9,258	12.9
22	32,730	26,655	44.9
23	15,126	7,337	32.7
24	7,081	4,840	40.6
25	1,253	948	43.1

Table 2: Descriptive Statistics

	Individuals	Percent with No Consumption	Average Expenditures, \$	Rates of Consumption, %
All	100,882	29.3	43.30	1.67
Loses Coverage	17,001	54.9	86.46	1.09
Does Not Lose	83,881	24.1	38.61	1.73
Loses in January	1,045	85.6	44.19	0.82
February	1,007	77.0	23.70	0.83
March	1,769	66.6	59.64	1.09
April	1,127	65.8	5.13	0.95
May	1,179	57.5	25.13	1.06
June	2,142	54.3	21.27	1.17
July	1,443	51.8	76.93	1.05
August	1,550	46.9	21.19	1.05
September	2,641	41.8	117.31	1.24
October	2,036	44.2	14.03	0.99
November	1,062	40.0	387.13	1.10

Table 3: Probability of Consumption at End of Coverage
 Dependent Variable: Percent of observations consuming care in the
 given category of health care

	(1)	(2)	(3)	(4)
	Total	Drug	Outpatient	Inpatient
Mean	1.18	0.59	0.72	0.01
Week 0	-0.159 (0.066) 0.04	-0.060 (0.038) 0.14	-0.128 (0.046) 0.02	-0.004 (0.003) 0.19
Week 1	-0.075 (0.054) 0.20	-0.031 (0.031) 0.34	-0.049 (0.041) 0.26	0.000 (0.002) 0.97
Week 2	-0.084 (0.056) 0.17	-0.042 (0.032) 0.21	-0.071 (0.037) 0.08	0.000 (0.003) 0.86
Week 3	-0.060 (0.058) 0.32	-0.011 (0.033) 0.75	-0.055 (0.040) 0.19	-0.002 (0.003) 0.41
Week 4	0.010 (0.043) 0.81	0.011 (0.025) 0.66	0.002 (0.032) 0.95	-0.001 (0.002) 0.60
Week 5	0.010 (0.056) 0.86	0.013 (0.027) 0.65	-0.005 (0.045) 0.92	-0.004 (0.004) 0.25
Week 6	0.060 (0.063) 0.37	0.003 (0.038) 0.94	0.071 (0.044) 0.13	0.004 (0.004) 0.38
Week 7	0.034 (0.053) 0.53	0.008 (0.021) 0.70	0.034 (0.047) 0.49	-0.001 (0.003) 0.70
Week 8	0.018 (0.046) 0.71	-0.008 (0.025) 0.74	0.018 (0.035) 0.62	0.002 (0.002) 0.51
Week 9	-0.026 (0.064) 0.70	-0.011 (0.032) 0.73	-0.018 (0.051) 0.73	-0.005 (0.002) 0.09
R ²	0.208	0.272	0.142	0.037

Note: The sample includes 2,363 daily cell means for individuals losing coverage on the same day. Standard errors in parantheses are robust to auto-correlation between observations ending coverage on the same day. All regressions include fixed effects for the week and last day of coverage (not shown).

Table 4: Consumption in Dollars at End of Coverage
 Dependent Variable: Mean expenditures for the given category of health care
 in dollars

	(1)	(2)	(3)	(4)
	Total	Drug	Outpatient	Inpatient
Mean	92.06	0.37	1.81	89.88
Week 0	-151.333 (77.781) 0.08	-0.038 (0.026) 0.17	-0.446 (0.187) 0.04	-150.849 (77.738) 0.08
Week 1	-79.227 (90.865) 0.40	-0.017 (0.031) 0.59	-0.294 (0.158) 0.09	-78.916 (90.821) 0.40
Week 2	-203.802 (157.968) 0.22	-0.037 (0.029) 0.22	-0.132 (0.297) 0.66	-203.633 (157.943) 0.22
Week 3	-293.832 (225.983) 0.22	-0.012 (0.036) 0.74	-0.270 (0.218) 0.24	-293.549 (225.985) 0.22
Week 4	-69.259 (33.974) 0.07	-0.016 (0.021) 0.47	0.407 (0.360) 0.28	-69.650 (33.880) 0.06
Week 5	-138.433 (77.730) 0.10	0.002 (0.025) 0.94	-0.190 (0.343) 0.59	-138.246 (77.679) 0.10
Week 6	-117.689 (83.594) 0.19	-0.020 (0.029) 0.51	0.694 (0.501) 0.19	-118.364 (83.459) 0.18
Week 7	-360.266 (362.916) 0.34	0.025 (0.025) 0.35	0.174 (0.243) 0.49	-360.465 (362.864) 0.34
Week 8	-57.019 (34.147) 0.12	-0.033 (0.044) 0.47	0.076 (0.173) 0.67	-57.062 (34.144) 0.12
Week 9	-137.153 (74.048) 0.09	-0.009 (0.026) 0.74	-0.078 (0.256) 0.77	-137.066 (74.102) 0.09
R ²	0.032	0.24	0.042	0.032

Note: The sample includes 2,363 daily cell means for individuals losing coverage on the same day. Standard errors in parantheses are robust to auto-correlation between observations ending coverage on the same day. All regressions include fixed effects for the week and last day of coverage (not shown).

Table 5: Robustness Checks

Dependent Variable: Percent of observations consuming any care, for the given sub-population

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Consumed in January	Low College Probability	High College Probability	Males Only	Females Only
Mean	1.18	3.23	1.62	1.70	0.81	1.61
Week 0	-0.159 (0.066) 0.04	-0.410 (0.262) 0.18	-0.170 (0.044) 0.00	-0.157 (0.077) 0.07	-0.114 (0.033) 0.01	-0.207 (0.113) 0.09
Week 1	-0.075 (0.054) 0.20	-0.506 (0.167) 0.03	-0.095 (0.057) 0.12	-0.069 (0.064) 0.31	-0.024 (0.040) 0.55	-0.134 (0.088) 0.15
Week 2	-0.084 (0.056) 0.17	-0.253 (0.174) 0.21	-0.095 (0.046) 0.06	-0.081 (0.062) 0.22	-0.067 (0.029) 0.04	-0.098 (0.104) 0.37
Week 3	-0.060 (0.058) 0.32	-0.522 (0.128) 0.01	-0.070 (0.071) 0.35	-0.060 (0.066) 0.38	-0.057 (0.044) 0.22	-0.060 (0.083) 0.48
Week 4	0.010 (0.043) 0.81	-0.010 (0.119) 0.94	0.010 (0.046) 0.84	0.011 (0.051) 0.84	0.014 (0.040) 0.73	0.015 (0.086) 0.87
Week 5	0.010 (0.056) 0.86	-0.243 (0.156) 0.18	0.033 (0.085) 0.70	0.005 (0.064) 0.94	-0.035 (0.044) 0.44	0.069 (0.089) 0.46
Week 6	0.060 (0.063) 0.37	0.065 (0.205) 0.76	0.155 (0.051) 0.01	0.041 (0.070) 0.57	0.060 (0.048) 0.24	0.064 (0.093) 0.50
Week 7	0.034 (0.053) 0.53	-0.287 (0.165) 0.14	0.031 (0.123) 0.80	0.037 (0.049) 0.47	0.008 (0.035) 0.81	0.074 (0.082) 0.39
Week 8	0.018 (0.046) 0.71	-0.055 (0.129) 0.69	-0.082 (0.058) 0.19	0.034 (0.058) 0.57	0.014 (0.050) 0.78	0.031 (0.083) 0.72
Week 9	-0.026 (0.064) 0.70	0.037 (0.219) 0.87	0.065 (0.110) 0.56	-0.043 (0.067) 0.53	-0.017 (0.059) 0.78	-0.032 (0.087) 0.72
R ²	0.208	0.302	0.134	0.209	0.21	0.151
N	2,363	1731	2363	2363	2363	2363

Note: The sample includes the given number of daily cell means for individuals losing coverage on the same day. Standard errors in parentheses are robust to auto-correlation between observations ending coverage on the same day. All regressions include fixed effects for the week and last day of coverage (not shown).