Claim Timing and Ex Post Adverse Selection: Evidence from Dental "Insurance"*

Marika Cabral†

March 30, 2012

Abstract

A large fraction of health care treatments are not urgent and may be delayed if patients so choose. Because insurance coverage is typically determined by the treatment date, individuals may have incentives to strategically delay treatments to minimize out-of-pocket costs. The strategic delay of treatment—a particular form of moral hazard—can be an important source of subsequent adverse selection, in which ex ante identical individuals select insurance coverage based on their differing accumulation of previously delayed treatments. This paper analyzes dental treatments and insurance, with the goal of understanding the insurance market for dental care and also revealing lessons that apply to insurance markets more broadly. Using rich claim-level data from a large firm, I present several simple tests of the hypothesis that people strategically delay dental treatments and adversely select insurance coverage. These tests provide the key identification and motivation for a structural model which I develop to explicitly link the endogenous delay of treatment to the adverse selection it causes, allowing me to evaluate the relative importance of this source of adverse selection as compared to traditional adverse risk selection. My analysis shows that the strategic delay of treatment and the associated adverse selection can explain why so few people have dental coverage in the US and why typical dental "insurance" contracts provide so little insurance. More generally, my results suggest that features such as open-enrollment periods and contracting on pre-existing conditions may overcome market unraveling in insurance contexts where the timing of risk is not contractible.

*I am especially grateful to Tim Bresnahan, Liran Einav, and Caroline Hoxby for their detailed feedback and invaluable guidance. I am also very grateful to Mark Cullen for his helpful advice and for his support in obtaining data for this project. In addition, I would like to thank Brenda Barlek and Martin Slade for their assistance in acquiring data and explaining the environment. For their useful comments and feedback, I thank Michael Boskin, Gopi Shah Goda, Matthew Harding, Jakub Kastl, Amanda Kowalski, Jonathan Levin, Neale Mahoney, Isabelle Sin, William Gui Woolston, and Ali Yurukoglu. Helpful feedback was provided by seminar participants at Brown, Columbia, RAND, STATA Texas Empirical Micro Conference, Stanford, UBC, UC Berkeley Haas School of Business, UC Irvine, University of Pittsburgh, University of Rochester Simon School of Business, USC, UT Austin, Wharton, and Yale School of Management. Support from the National Science Foundation is gratefully acknowledged.

†Department of Economics, UT Austin, 2225 Speedway Stop C3100, Austin TX 78712-1690, marika.cabral@austin.utexas.edu.

In most non-emergency medicine, there is a time lapse between the recognition that a particular treatment is necessary and the actual treatment itself. This time lapse can often be controlled by the patient, so that treatment is delayed if the patient wishes it to be. Because insurance coverage is typically determined by insurance status on the treatment date, such control over the timing of treatment can generate substantial problems for insurance. If people can delay a treatment (once they know it will occur) just long enough to buy more insurance in anticipation of it, severe adverse selection may result. In the extreme case where all treatments can be easily delayed, individuals can opt out of insurance, accumulate untreated problems, buy generous coverage and get treatment for the entire stock of problems they have accumulated, and then opt out of insurance again; in such cases, these forces can lead insurance markets to completely unravel. Unlike typical adverse selection, which is driven by heterogeneity in ex ante risk types, this "ex post adverse selection" may arise even in a homogeneous population, as individuals become differentiated over time by the amount of treatment they have postponed. Ex post adverse selection has different policy implications than traditional moral hazard and traditional adverse selection. While the insurance coverage period is of little consequence in the context of traditional adverse selection (or moral hazard), the best policy response to ex post adverse selection may include waiting periods, open-enrollment periods, and eligibility restrictions based on prior insurance status.

Health care treatments span a spectrum of urgency. While treating a heart attack is extremely costly to delay, most orthopedic surgery can be delayed easily for quite some time. Because the treatments generally covered by medical insurance range from extremely urgent to very delayable, medical insurance is a mixture of a market affected by ex post adverse selection and a market unaffected by it. This makes identifying the effects of strategic timing challenging. In contrast, only a very small share of dental treatments are extremely costly to delay,[1] making the insurance market for dental care a natural venue for investigating how strategic timing affects insurance coverage.[2] In this paper, I analyze dental treatments and insurance, with the goal of understanding not only insurance in the \$100 billion dental market[3] but also revealing lessons that apply to insurance markets more broadly. More generally, these timing and selection incentives may play a role in any insurance setting where (1) individuals can control the timing of the insured financial cost after the realization of an "event" that necessitates this cost and (2) the date of the financial cost, rather than the date of the "event," is used to determine coverage. In other words, these incentives are relevant in contexts where the timing of claims is manipulable and the timing of uncertain events is not contractible. These two conditions are often met in the cases of medical and dental insurance; thus, in these contexts, there

[1] A number of articles in dental journals describe the flexibility in timing dental care (see, e.g., Guay 2006, Jeffcoat 2004).

[2] The historical evolution of medical insurance in the United States has excluded dental care from coverage. In many ways, the exclusion of health care related to the mouth from general health insurance seems arbitrary. This paper illustrates that a separate market for dental insurance is not sustainable because of the high degree of delayability. The welfare effects of altering the bundling services covered by insurance (for example, bundling dental with the rest of health care) are ambiguous. The optimal bundling of services for insurance is a potentially promising area for future research.

[3] Total US dental spending in 2008 was \$100 billion according to Palmer (2009).

is a relatively large scope for strategic timing.[4]

Dental risk is largely uninsured in the United States. Only about 40% of individuals in the US have any dental coverage,[5] and those with coverage actually have little insurance against dental risk. The typical dental insurance policy provides very incomplete coverage owing to a low annual maximum benefit, which is on average $1,100. Above this maximum benefit, insured individuals must pay the full cost of services. Although dental care can involve considerable uncertainty and financial cost, available insurance policies tend to offer no coverage for large, urgent dental expenditures. Perhaps because the available policies provide so little insurance, almost no one takes them up except through an employer. That is, if it were not for the tax subsidy that allows people to pay for dental premiums with pre-tax dollars when they enroll in an employer-sponsored plan (as opposed to post-tax dollars when they pay for care out of pocket), dental "insurance" might not exist at all. From the perspective of a researcher, it is fortunate that this tax subsidy exists as companies that offer these subsidized benefits collect dental spending data in order to review claims. There has been very little prior research explaining the non-existence of insurance markets, perhaps because of the difficulty in obtaining data related to missing markets. One of the contributions of this paper is that I can analyze the largely missing market for dental insurance because complete dental spending data are collected by employers.[6]

In this paper, I use rich and complete claim-level data from one large firm, Alcoa Inc., to investigate the prevalence of strategic treatment delay and its consequences. The data are particularly useful for at least four reasons. First, the firm offers two vertically differentiated dental plans, where the primary difference between the plans is the size of the maximum benefit, $1,000 or $2,000. This feature allows me to look for evidence of selection across these options. Second, because the firm's less generous dental insurance plan is free for employees, no one opts out of dental benefits. Therefore, the data contain all dental claims for all employees for all years they are with the company. Third, the data include claims for treatments that were not reimbursed because the costs were above the annual maximum benefit. Thus, the data are not censored. Fourth, during the data period (2004-2007), the firm changed the coverage period twice (going from one-year coverage to two-year coverage and then back to one-year coverage). The changes in the coverage period are very useful for identification, as will become clear below.

Analyzing the strategic delay of treatment and the subsequent incentive to buy more insurance coverage is inherently challenging. The decision about when to time a treatment, conditional on the recognition that treatment is necessary, is a complex, dynamic decision. A person needs to take into account his inventory of untreated events (if any), the insurance choices that will be available

[4] In the case of medical and dental insurance, individuals can often delay costly treatment after the realization of an event (a health problem), and coverage is typically determined by insurance enrollment on the treatment date. In contrast, strategic timing plays little role in auto insurance where the date of a typical event, an auto accident, is generally contractible so insurers need not rely on the date of the associated auto repairs to determine coverage.

[5] Calculation based on self-reported insurance status in the Medical Expenditure Panel Survey, 2007.

[6] There are many risks for which insurance is not available. However, typically it is difficult to obtain data on these risks so researchers have made little traction in studying these missing markets. Because employers collect data through administering incomplete dental coverage, I am able to study a the missing market for comprehensive dental insurance in this paper.

to him in the future, and the expected premium of these future insurance options. In addition, we (the econometricians) do not observe when an event that requires treatment is realized or the costs of delaying this treatment. Given this challenging environment, I adopt two complementary analytic approaches. My first approach is a series of reduced-form tests for the patterns in the data that theory predicts will exist when strategic timing and ex post adverse selection are important. The advantage of this approach is that I find clear and transparent evidence of both strategic claim delay and dynamic asymmetric information. This evidence also makes us aware of the variation that ultimately identifies the more structural analysis, which is my second approach.

For example, in my first analytic approach, I look for a spike in dental treatment (and spending) in the beginning of the calendar year among only those people who likely had an incentive to delay treatment because they were likely to have exceeded the previous year's maximum benefit. The size of this spike should depend on whether the adjacent years were during the two-year coverage period (the annual maximum resets to zero but a person cannot choose a more generous policy) or whether people could change coverage between the adjacent years (the annual maximum resets and a person can choose a more generous policy). In addition, I inspect the data for evidence of dynamic asymmetric information. To do this, I look for a positive correlation between claims and choice of more generous coverage, a classic symptom of asymmetric information (Chiappori and Salanie, 2000). My analysis shows there is strong evidence of asymmetric information in this setting and that much of this asymmetric information operates within-household, over time, consistent with ex post adverse selection.

The tests employed in my first approach provide clear evidence of strategic treatment delay and dynamic asymmetric information. However, this evidence alone does not allow me to analyze counterfactual policies, such as contracting on pre-existing conditions and open-enrollment periods. This motivates my second analytic approach: a structural model that explicitly links strategic claim delay to the adverse selection it can cause. The model I develop allows me to quantify the relative importance of ex post adverse selection (as opposed to traditional adverse risk selection) and to explore counterfactual policies. In the model, individuals realize dental events and then decide how much of the associated treatment to delay until the following year, knowing they will receive a stochastic draw of events in the following year that depends on their risk type. The model makes explicit the cost of delay (e.g., ongoing pain or cognitive costs). The key model primitives I estimate are unobserved heterogeneity in delay costs and unobserved heterogeneity in dental risk. Intuitively, the distribution of dental risk is identified by the distribution of total claims in the data (without regard for year-to-year timing), and the distribution of delay costs is identified by how individuals close to their maximum benefit allocate claims between adjacent years. Note that this latter source of identification is one of the key correlations examined in my first, more descriptive approach.

The model estimates suggest that a considerable share of treatments can be delayed. Approximately 48% of individuals delay claims from one year to the next when they have incentives to do so. Using the model estimates, I extrapolate beyond the firm's benefits to investigate the missing market for comprehensive dental insurance. My counterfactual analysis illustrates that an unsubsidized market for annual comprehensive dental insurance would unravel if insurers could not price both pre-existing events (the source of ex post adverse selection) and ex ante risk information (the source of traditional adverse selection). This suggests that even if insurers were able to perfectly identify and price ex ante risk information, the non-contractibility of pre-existing events would still make an annual comprehensive dental insurance product unviable. In addition, I evaluate the effectiveness of policies that may address this adverse selection. In equilibrium, fewer than 5% would enroll in comprehensive dental insurance without substantial reforms relative to typical insurance markets such as: pricing more risk information, significantly restricting insurance choice frequency (thereby increasing coverage periods), and expanding premium subsidies. Overall, I show that the strategic delay of treatment and ex post adverse selection is one explanation for why so few people have any dental coverage in the US and why available policies offer so little insurance. More generally, my results suggest that strategic delay and the associated adverse selection may motivate contract features seen in many insurance settings and may explain why some risks are difficult to insure. In particular, the results suggest that features such as open-enrollment periods and contracting on pre-existing conditions may overcome market unraveling in insurance contexts where the timing of risk is not contractible and claims are delayable.

The remainder of the paper proceeds as follows. Sections 1 and 2 describe, respectively, the related literature and the data. Section 3 discusses the theoretical incentives for delaying treatment and selecting coverage ex post in the context of the studied firm's dental benefits. Section 4 presents evidence that people behave according to these incentives. Section 5 describes the structural model—its setup, identification, estimation, and results. In Section 6, I analyze counterfactual policies. Lastly, I conclude in Section 7 by summarizing my findings and describing the relevance of the results for broader health insurance markets.

1 Related Literature

Building on the seminal work of Akerlof (1970) and Rothschild and Stiglitz (1976), a growing body of recent literature has focused on identifying and quantifying the impact of asymmetric information in insurance markets. This empirical literature has evolved from theoretically inspired tests for asymmetric information to efforts aimed at quantifying the welfare effects of asymmetric information. Much of this literature builds on the work of Chiappori and Salanie (2000), who outline a robust set of tests for asymmetric information. The basic idea behind these tests is that, in the presence of asymmetric information, one should observe a positive correlation between the generosity of the chosen insurance coverage and claim realization conditional on the information priced by the insurer. The results of many subsequent papers that use some version of this positive correlation test have been mixed: some papers find little evidence of asymmetric information in particular markets (e.g., Chiappori & Salanie 2000, Cardon & Hendel 2001), while some studies find evidence of asymmetric information in other markets (e.g., Finkelstein & Poterba 2004, Cohen 2005). In addition to the standard explanations for a positive correlation in claims and coverage (traditional adverse selection and traditional moral hazard), the strategic timing of claims and the associated ex post adverse selection studied presently may lead to such positive correlation. In my first empirical approach described in Section 4, I look for correlations indicative of asymmetric information using a similar approach.

A number of recent papers go beyond testing for the presence of asymmetric information and seek to quantify the effects of asymmetric information in insurance markets. For a comprehensive review of this literature, see Einav, Finkelstein, and Levin (2010). An early work in this spirit is by Cardon and Hendel (2001), who model not only adverse selection based on private risk information but also moral hazard in the utilization of medical services. Cohen and Einav (2007) develop and estimate a model that uses deductible choice in automobile insurance to estimate the joint distribution of risk aversion and claim risk. Using a similar modeling approach, Einav, Finkelstein, and Schrimpf (2010) estimate a model of annuity selection which they use to investigate the welfare effect of mandates in the UK annuity market. Bundorf, Levin, and Mahoney (2010) use a more traditional discrete choice framework to analyze the effect of offering HMO- and PPO-type health insurance plans on pricing and welfare. Einav, Finkelstein, and Cullen (2010) develop a simple framework that can be used to recover the willingness-to-pay for insurance and to evaluate the welfare cost of asymmetric information. A number of other papers (e.g., Carlin & Town 2009, Sydnor 2009, Handel 2010, Bajari et. al. 2010 and Lustig 2010) attempt to quantify private information in various insurance settings. Most of the papers in this literature focus on modeling asymmetric information in static insurance settings. The empirical model presented in Section 5 builds on these previous works by modeling asymmetric information that is inherently dynamic.

This paper is also related to the literature on dynamic inefficiencies in insurance, and is related in particular to the work on reclassification risk and the difficulty of insuring long-term risk (e.g., Cochrane 1995, Hendel & Lizzeri 2003, Finkelstein, McGarry & Sufi 2005). These papers focus on inefficiencies that can arise in insurance markets even when insurers and consumers have symmetric information. Such dynamic inefficiencies can arise in the absence of enforceable lifetime contracts because individuals are typically not fully insured against becoming a bad risk and being reclassified into a higher risk group associated with higher insurance premiums. While the inefficiency I study in this paper also exists in the absence of enforceable lifetime contracts, the source of the inefficiency here is the non-contractibility of underlying events (and thus the reliance on the timing of treatment to determine insurance coverage). Thus, in contrast to the literature on reclassification risk, the inefficiency studied in this paper stems from the inherent asymmetric information between insurers and consumers that arises when insurers cannot observe the timing of events and individuals can both delay costs and re-evaluate insurance decisions.

Other related papers have estimated the response of health care utilization to insurance coverage. The gold standard in this literature are the studies based on the Rand Health Insurance Experiment, a large-scale experiment from the mid-1970s in which households were randomly assigned to health insurance plans with differing levels of benefits. Using data from this experiment, Manning et al. (1987) estimate the price elasticity for overall health care spending to be around -0.2. Earlier papers by Arrow (1975) and Metcalf (1973) point out the difficulty of inferring long-run policy responses from such temporary experiments due to transitory demand (potentially induced by the sort of strategic timing I study in this paper). Indeed, the Rand insurance experiment revealed some data that are consistent with the strategic postponement of care. Manning et al. (1986) show that the responsiveness of spending to insurance coverage is larger in the first year of the experiment than in subsequent years, and this elevated first-year responsiveness is much more pronounced in dental spending than in other health care spending. Using more recent non-experimental data, Kowalski (2010) identifies the price elasticity of medical care utilization using a novel instrumental variable approach. Using the injury of a family member as an exogenous shift in the marginal price for health care, the author estimates the price elasticity of utilization to be around -2.3, which is an order of magnitude higher than that found in the Rand experiment. In addition to the many theoretical and econometric reasons for the difference in these estimates discussed by Kowalski (2010), the strategic timing of care may be an alternative explanation for some of this discrepancy. Card, Dobkin, and Maestas (2008) reveal some evidence consistent with the strategic delay of medical treatment in anticipation of better insurance coverage. The authors examine health care utilization and access to care near 65 years of age, the Medicare eligibility threshold. Their analysis reveals that the number of individuals who report delaying treatment decreases at this threshold, and there seems to be a modest increase in physician visits at the age of eligibility, with larger gains for groups with lower rates just before age 65. Examining the Flexible Spending Account (FSA) contributions and claims of a medium-sized insurer, Cardon and Showalter (2001) find some suggestive evidence that people delay dental spending. The authors find that approximately 40% of FSA claims for dental spending reimbursement were for dental spending that took place in January. Overall, the authors find that FSA funds tend to be used to reimburse highly predictable spending, and the majority of employees exhaust their FSA balances well before the end of the calendar year.

The particular ex post channel for adverse selection studied in this paper arises because individuals delay claims and subsequently switch insurance coverage. Though there have been many studies on switching behavior in related contexts, most of these studies focus on switching induced by forces external to individuals: changes in prices (e.g., Buchmueller & Feldstein 1997, Cutler & Reber 1998, Handel 2010), changes in defaults (e.g., Madrian & Shea 2001, Goda & Manchester 2010), and changes in the set of options (e.g., Strombom, Buchmueller & Feldstein 2002, Abaluck & Gruber 2009). More closely related to this paper, a few studies have described the correlation between health insurance switching and changes in health status. Cutler, Lincoln, and Zeckhauser (2010) document three types of health-related transitions between health insurance plans: (1) "adverse selection," the movement of the less healthy to more generous plans, (2) "adverse retention," the tendency for people to stay where they are when they get sick, and (3) "aging in place," enrollees' inertia in plan choice, leading plans with older enrollees to increase in relative cost over time. Using data from Massachusetts, the authors show that adverse selection and aging in place are both quantitatively important explanations for transitions between health insurance contracts. Tchernis et al. (2006) find evidence that people who switch to more generous health insurance plans do not have particularly high levels of spending before the switch but increase medical expenditures substantially after switching. The authors find that much of this spending increase is for mental health services. They suggest these findings are consistent with the strategic delay of mental health treatment in anticipation of switching plans, and the authors discuss how such strategic delay could exacerbate adverse selection. My paper builds on this work in many ways. I formalize the intuition of strategic claim delay and ex post adverse selection, and estimate the extent of this behavior in a more controlled context. Using the model estimated in Section 5, I investigate the relative importance of strategic timing in generating adverse selection compared to more traditional sources of adverse selection. In addition, I estimate the effect of policies that may be particularly important in these contexts such as insurance enrollment frequency restrictions.

2 Data

I use rich, claim-level data from a self-insured, multinational manufacturing company that employs approximately 48,000 individuals in 40 states across the US. The company offers dental benefits to employees and their dependents, and covers approximately 110,000 individuals through dental benefits annually. The data contain claim-level information regarding dental and medical spending as well as insurance coverage choices from 2004 to 2007. In addition, the data contain basic demographic information for employees, including wage, sex, age, and job tenure. The data include claims for all employees and insured dependents.

Each claim contains information on the total claim cost, out-of-pocket expenses, insurance payment, date of service, and procedure codes, which vary in level of specificity. While the data contain financial information for each separately billed procedure, I aggregate this information to the individual-day level, and "claims" in the remainder of the paper will refer to the individual's total billed procedures within one day.[7] A crucial feature of the claims data is that it contains information on *all* dental spending for insured individuals. All claims submitted to the insurance administrator are reported in the data, including unpaid claims after the individual has exhausted his annual benefits.

2.1 Definition of Baseline Samples

Dental insurance options vary within the company by employee benefit groups. Employees are divided into benefit groups that reflect the firm's subsidiary business model, as well as employee occupation, job location, and union membership. The analysis in this paper focuses on the most common dental benefit menu that was offered to approximately 70% of employees over the observed years. While the company introduced this menu in 2004, the company rolled out this dental benefit menu over a number of years because of staggered union contract expirations.

Employee demographic information is described in Table 1 for each of the samples analyzed in this paper. The first column describes all employees who were ever observed in the data over the available years, 2004-2007. The second column describes the employees who have the relevant benefit menu offered to them at some point during the observed years; these employees along with their associated

[7] There are typically multiple procedure codes billed for a dental visit. Some dental visits may necessarily involve multiple procedure codes because of the complexity of treatment, while other visits may contain multiple procedures by the patient's choice. While ideally a claim would be defined as a related collection of procedure codes that cannot be unbundled by patient choice, the data do not contain the information necessary to make this distinction. In the remainder of the paper, the word "claim" is used to describe the total procedures claimed on a particular date for an individual.

dependents will be referred to as the “baseline sample.” In Section 4, the baseline sample is used to look for evidence of claim timing and adverse selection. The third column summarizes the employees in the sample used to estimate the structural model. This sample is restricted to employees and dependents associated with employees who were offered the relevant dental benefits in the two years used in the estimation of the model[8] and this sample will be referred to as the “restricted sample.”

From inspecting Table 1, one can see that the median employee tenure is about seven years. The majority of employees are male, and about 40% of employees live in rural areas. The median wage is around \$42,000, and the median employee age is 44 years. Approximately 70% of employees choose to enroll dependents in dental insurance. The individuals in the restricted sample look a bit different from those in the overall company population; fewer of those in the restricted sample are unionized[9], their earnings are slightly higher on average, and fewer of them opt for employee-only coverage. One can compare the employee demographic information against a representative sample of employed individuals with dental insurance in the US. The last column in Table 1 displays some descriptive statistics for the sample of individuals in the Medical Expenditure Panel Survey (2007) who are continuously employed and enrolled in dental insurance throughout 2007.[10] The median age and mean wage look similar in the company and the overall employee population. Compared to the overall employee population, a much larger fraction of the company employees are male and unionized. Dental spending of individuals in the baseline sample and the overall US population are compared in detail in Appendix A.

2.2 Description of Plan Details

Table 2 describes the dental insurance benefits of Plan L and Plan H, the two plans available on the relevant benefit menu. This table reports the percentage of dental expenditures that the company will pay below the annual maximum benefit by dental claim category: basic care, major care, oral surgery, and preventive care. Plan H has a \$2,000 annual individual maximum benefit, while Plan L has a \$1,000 annual maximum benefit per covered individual. Once individuals reach this annual maximum benefit, the insurer does not reimburse for subsequent dental treatment. In addition, there are some small differences in the coinsurance rate below the maximum benefit for some dental care.

The company varied the length of commitment to this annual insurance coverage over the time period studied. In particular, the company “locked” employee dental insurance decisions for two years, 2005-2006, and later reverted to annual dental coverage decisions. Coverage decisions are made during November of the year preceding the calendar year for which the coverage first applies. By November, it is likely that individuals will know almost all their dental problems for the current calendar year and will generally know whether they have delayed claims to the next year. It is

[8]Conditional on being an employee in a given year, 82.5% of employees remain with the company in the following year; other employees part with the company due to termination, suspension, and retirement.

[9]Those in the restricted sample are offered the relevant dental benefits in both 2005 and 2006. Many of the unionized employees in the company do not meet this criteria because the company only switched them to the relevant benefit menu after 2005 due to union contract expiration dates.

[10]This is not a perfect comparison group as the individuals in this sample may have obtained dental insurance through sources other than their own employer (for example, their spouses’ employer).

important to note that even during the years with locked insurance coverage, the insurance terms still included an annual maximum benefit that applied in each calendar year. This means that, even during a locked period, individuals have the incentive to delay claims from one year to the next if they require dental treatment that puts them at the individual annual maximum benefit threshold during the first year of the locked period.

Similar to the majority of firms that offer dental insurance to employees, the company subsidizes this insurance in the sense that premiums are lower than the average cost of insured individuals. Plan L is available to all employees and dependents at no cost, while Plan H is available for a premium. It is convenient that Plan L coverage is free to employees as this means there is universal coverage for dental care among employees, so all employee dental usage is recorded by the company and available in the data. The premium for Plan H depends on the chosen coverage tier: employee-only coverage, employee plus family coverage, employee plus spouse coverage, or employee plus children coverage[11]. In addition, the Plan H annual premium varies across benefit groups and over time. In 2004, the average Plan H premium for family coverage was just under \$150 while the average premium for Plan H single coverage was around \$50. In 2005, average Plan H premiums increased to \$200 and \$65 for family and employee-only coverage respectively. Plan H premiums remained roughly constant for the remainder of the sample period. Figure A1 in Appendix A displays the average premium for Plan H by coverage tier over the years.

Employees that select Plan H coverage pay the associated premium with pre-tax income. Additionally, some out-of-pocket dental expenditures may also be paid with pre-tax income saved in tax-advantaged accounts offered by the company.[12] Because the data does not contain household income to infer income tax brackets nor does the data include information on which out-of-pocket expenses were paid with funds from tax-advantaged accounts, premiums and out-of-pocket expenses are treated symmetrically as post-tax expenditures in this paper. As will become clear later, the estimation of the structural model is not sensitive to the tax treatment of premiums.[13]

The reimbursement of dental spending under the two plans varies somewhat with the category of care. Examples of these categories are given in the employee dental plan information brochure: basic care (fillings, root canal therapy), major care (bridgework, dentures), preventive care (exams, cleanings, emergency pain treatments), and oral surgery (removal of impacted teeth). Table 2 displays the breakdown of claims and spending for the baseline sample by inferred category of care.[14]

[11]Because the Plan H premium varies with the coverage tier, it is possible that some dependents go uninsured in some years. The selective enrollment of dependents is another potential avenue for adverse selection (in addition to plan choice) and is dealt with in detail in Section 4.

[12]Employees may use Flexible Spending Account funds or Health Savings Account funds to pay for out-of-pocket dental or medical expenditures.

[13]The estimation is insensitive to the tax treatment of premiums because (1) premiums remain stable over the time period used to estimate the model, and (2) the model is estimated during a period in which coverage decisions were locked. The only avenue through which the tax treatment can affect the model is through the calibration of the risk aversion parameter. Robustness analysis in Appendix C reveals that the estimates are not too sensitive to this calibration.

[14]Claim categories are inferred by combining procedure codes and the claim reimbursement information. The average out-of-pocket spending to total spending ratio is calculated for each procedure code, and these codes are then classified into care categories. This process left less than 5% of procedures with unclassifiable codes. Claims with these codes are omitted from the statistics on the percentage of procedures and spending by category.

Approximately 57% of claims are for preventive care, while 42% claims are for basic care. However, 63% of dental spending is basic care spending, while 32% is preventive care spending. The remaining spending and claims are for major care and oral surgery. For simplicity, the rest of the paper will abstract from these categories of care.[15] Individual out-of-pocket spending on Plan j as a function of total spending is calculated using the average coinsurance rate, γ_j, defined as the average percentage of spending the company reimburses across the categories of care below the maximum benefit.[16] I allow this average coinsurance to vary with age in a categorical manner to capture the fact that the types of treatments done by the middle-aged differ from those done by children. Averaging across age categories in the baseline sample, the average coinsurance rate for Plan L is 84.2% below the maximum benefit of $1,000, and the average coinsurance for Plan H is 87% below the maximum benefit of $2,000.

Figure 1 plots individual out-of-pocket expenditures (excluding premiums) as a function of total individual expenditures using the baseline sample plan average coinsurance rates. For Plan j, one can write this function as follows:

$$\mathrm{OOP}_j(x) = x - \min(\gamma_j x, b_j). \tag{1}$$

Because this calculation excludes premiums and the two plans are vertically differentiated, in the figure the out-of-pocket expenditures on Plan L are greater than the out-of-pocket expenditures on Plan H for any given level of total expenditures. Inspecting the figure, one can see that the main difference between the plans is the difference in annual maximum benefits. One can calculate the amount of individual dental spending it would take to exhaust the annual maximum benefit of each plan (the spending levels that correspond to the location of the kink points in Figure 1). An individual facing the average coinsurance rates above would exhaust the annual maximum benefit by spending around $1,188 on Plan L or $2,299 on Plan H. To get a sense of how much dental spending one would need to be better off ex post having Plan H coverage, one can compare the Plan H premium to the out-of-pocket spending differences displayed in Figure 1. Given the employee-only premium of $65, a single employee would be better off *ex post* under Plan H if he had more than approximately $1,225 of dental spending.

2.3 Description of Dental Claims and Spending

Figure 2 panel a displays the distributions of annual individual dental expenditures and claim cost for the baseline sample side-by-side. While the majority of claims cost less than $200, the sparse right tail of claims reaches a few thousand dollars. The mean cost of a claim is $161 with a standard deviation of $224. The mean annual individual dental spending in the baseline sample is $257 (median $107), while approximately 38% of individuals have no dental expenditures in a given year. The right tail of dental expenditures is thin with the maximum observed annual spending around

[15] The empirical model abstracts from categories of care as detailed in this section. Some of tests for correlations in Section 4 and the appendix use actual out-of-pocket payments reported in the data, which are of course conditioned on spending by type of care. When actual out-of-pocket payments are used, it is noted in the table caption.

[16] This is basically a weighted average of the coinsurance rates for the different categories of care. Specifics of this calculation are described in Table 2.

$20,000. Figure 2 panel b displays the right tail of annual individual dental spending by plan. Inspecting the figure, one can see that the percent of individuals with dental spending falls sharply near the point at which individuals would exhaust the maximum benefit of each plan, around $1,188 for Plan L and $2,299 for Plan H. There also seems to be some bunching of individuals near the level of spending it would take to exhaust the annual maximum benefit of Plan L. Though Figure 2 is created using the baseline sample, Appendix A includes figures that demonstrate that the annual expenditure and claim cost distributions look very similar in the restricted sample.

Table 3 displays descriptive statistics by plan enrollment for the samples used in the following sections. The baseline sample includes 118,112 individuals across the years, while the restricted sample contains 46,271 individuals. Approximately 76% of household-years are enrolled in Plan H in the samples. A large fraction of individuals on both plans have zero expenditures in a given year. Thirty-eight percent of individuals on Plan H in the baseline sample have no dental expenditures in a given year despite paying a premium for Plan H coverage. The large fraction of individuals and families with no dental spending selecting Plan H indicates that there is some uncertainty in dental spending and many individuals value the available coverage of this dental uncertainty despite the low annual maximum benefit.[17] Individuals enrolled in Plan H have higher expenditures than those enrolled in Plan L: $70 more expenditures on average in the baseline sample, and $80 more expenditures in the restricted sample. In the baseline sample, 2.1% of individuals enrolled in Plan L reach the $1,000 annual maximum benefit, while 3% of Plan L enrollees in the restricted sample reach this maximum benefit. Less than 1% of individuals enrolled in Plan H reach its $2,000 maximum benefit in either sample. The timing incentives explored in the following section may explain why relatively few individuals have expenditures that reach or exceed the annual maximum benefit. Significantly more individuals get within $200 of exhausting the individual annual maximum benefit. Among those on Plan L, 6.5% in the baseline sample and 5.4% in the restricted sample get close to the maximum benefit in a given year, while among those on Plan H, 0.9% in the baseline sample and 1.2% in the restricted sample get close to the maximum benefit.

3 Claim Delay and Selection Incentives

I discuss two general incentives for delaying claims in the present context and then discuss how delaying claims may lead to the subsequent incentive to select more generous insurance coverage ex post. Before continuing, it is worth reiterating an important distinction for the following discussion and the remainder of the paper: an "event" is a problem that requires treatment (for example, a dental cavity), and a "claim" is the treatment of an event (for example, the associated filling).[18] After realizing an event, an individual may decide to strategically delay the associated treatment (claim).

Two insurance incentives may motivate individuals to strategically delay claims. First, per-period nonlinearities in insurance coverage, in this context the annual maximum benefit, can incentivize

[17] This demonstrated risk aversion can be seen more clearly through inspecting those who select employee-only coverage. Of those employees that select employee-only Plan H coverage, 43% have no claims in a given year.

[18] Since all spending is observed in the data, I make no distinction between treatment and claims in this discussion.

individuals to delay claims. When treatment costs exceed the annual maximum benefit threshold in a given year, an individual has an incentive to delay costs beyond this maximum until his benefits reset next January. In this way, the individual gains coverage for treatment that he would have paid completely out-of-pocket otherwise. Second, the opportunity to select more generous insurance coverage after the realization of events may motivate individuals to delay claims. In the environment studied presently, the opportunity to select more insurance coverage, to switch from Plan L to Plan H, may motivate individuals to delay more claims than they would otherwise delay if they were restricted to remain on Plan L. The intuition behind this is simple. An individual enrolled in Plan L who requires a lot of treatment beyond the maximum benefit, would have an incentive to delay some amount of treatment even if he was restricted to remain on Plan L, as in the period when the company locks coverage. However, if he had the opportunity to sign up for Plan H in the coming year, he may want to delay more treatment to take advantage of the higher Plan H maximum benefit. After delaying claims due to these incentives, individuals may have a subsequent incentive to select more generous insurance coverage. This subsequent ex post adverse selection incentive stems from the fact that an individual who has postponed treatment anticipates future treatment costs which generally increases his valuation of Plan H benefits. Appendix A contains a simple example to illustrate these incentives more explicitly.

In reality, many frictions may prevent individuals from delaying care or selecting more generous insurance ex post. For example, individuals may find it costly to delay claims because of pain or inconvenience associated with delaying treatment. In addition, cognitive limitations could inhibit both strategic claim delay and subsequent plan switching. Results described in Section 4 show that these frictions are not too large because there is evidence of both a substantial amount of strategic treatment delay in the data and evidence of asymmetric information associated with plan switching. Using a structural model, Section 5 estimates that the extent of strategic timing is substantial, and these estimates are then used in Section 6 to investigate various counterfactual policies.

4 Evidence of Incentivized Behavior

This section examines the data for evidence of behavior encouraged by these timing and selection incentives. First, I show that coverage choices and claim realizations are positively correlated, indicative of some asymmetric information in this context. Consistent with ex post adverse selection, the results suggest that much of the asymmetric information in this context operates dynamically within-household, over time. Second, direct evidence of strategic claim delay is presented. Taking advantage of the annual maximum benefit contract feature, I use several tests to show that incentivized individuals postpone treatment until just after the commencement of a new calendar year (at which point their benefits reset). This displacement of spending from the end of one calendar year to the beginning of the next lines up exactly with individuals' incentives to delay care in this environment.

4.1 Asymmetric Information

A central theoretical prediction in many models of asymmetric information is a positive correlation between coverage and claim realization (Chiappori & Salanie 2000). This positive correlation can arise for several reasons. Traditional moral hazard incentives may lead people with more comprehensive insurance to have more discretionary dental spending. People who are ex ante more risky may select more coverage, as in traditional adverse risk selection. Ex post adverse selection may lead to a positive correlation in claims and coverage over time as individuals sign up for more insurance in anticipation of treating delayed events. Following previous empirical studies (e.g., Chiappori & Salanie 2000, Finkelstein & Poterba 2004), I look for evidence of positive correlation in coverage choice and claim realization to test the joint hypothesis that there is some type of asymmetric information in this setting. To do this, I estimate the following equation:

$$Claims_{h,t} = \gamma + \alpha Choice_{h,t} + X_{h,t}\beta + \epsilon_{h,t}. \tag{2}$$

The basic idea behind this test is to look for a positive correlation between claim realization, $Claims_{h,t}$, and insurance coverage choice, $Choice_{h,t}$, conditional on the household information priced by the insurer, $X_{h,t}$. The test for positive correlation is then a test of the sign and significance of α. Some possible sources of asymmetric information may contribute to cross-sectional correlation, while other sources contribute to within-household correlation over time. To distinguish between the within- and across-household correlation in claims and coverage, the model is estimated both with and without household fixed effects taking advantage of the panel nature of the data.

To perform this test, it is important that the insurance options can be ranked by generosity, with $Choice_{h,t}$ indicating the generous option. Fixing the number of dependents covered by an employee, the company dental insurance plans, Plan L and Plan H, satisfy this condition. However, employees may select the number of family members to insure and have an incentive to select this number carefully due to premium variation based on the number of insured family members. This ability to select the coverage tier means that in reality there are more than two coverage options for most employees, and some of these options are not ranked vertically. For example, consider a typical employee with a spouse. The employee can insure both himself and his spouse under Plan L for free, insure only himself under Plan H for \$65, or insure himself and his spouse under Plan H for \$145. Notice that while the first two options cannot be ranked in terms of generosity of coverage, the third option is strictly more generous than the first two options. The maximum coverage available to any employee is to insure his entire family under Plan H.

In practice, two different definitions of $Choice_{h,t}$ are used in the analysis. In the first specification, I define $Choice_{h,t}$ such that it indicates when the employee has selected to insure all his "potential dependents" under Plan H in year t. An employee's "potential dependents" are defined as all dependents who are ever covered by the employee through either dental or medical insurance in year t. This definition is meant to approximate the true family structure of the employee to the extent that it is observed in the data in order to compare vertically ranked insurance options for the household. In the second specification, I restrict the sample to those households that do not seem to strategically select which dependents to insure and define $Choice_{h,t}$ as an indicator that household

h is enrolled in Plan H in year t. In this specification, the sample is limited to households that in each year insure the same dependents for medical and dental insurance.[19] The systematic selection of dependents to enroll in insurance is another avenue through which adverse selection can operate; the first definition of $Choice_{h,t}$ includes selection of this sort, while the second definition captures asymmetric information net of this dependent selection.

Let $Claims_{h,t}$ be the amount of money the company would have paid in claims for household h in year t, had the whole household been enrolled in Plan H in year t (regardless of the actual household enrollment).[20,21] In this context, this measure is more desirable than coarser measures of claims because it can capture the finer differences between the plans that may cause asymmetric information to be important.[22] The covariates $X_{h,t}$ are those exogenous household characteristics the company uses to price $Choice_{h,t}$. In this employer-provided insurance setting, very little information is priced. Because there is some variation in the premium menu employees receive based on occupation and location, I control for the premium menu in the regression above.[23]

Table 4, columns (1) and (3) display the OLS regression results without household fixed effects for the first and second definitions of $Choice_{h,t}$, respectively. The results show that there is a very significant positive correlation between claim realization and coverage choice. Insurer expenditures per household under the maximum coverage are on average \$267 higher in the first specification (and \$206 higher in the second specification) for households enrolled in the maximum coverage than for households enrolled in less coverage. However, these estimates confound sources of asymmetric information that would cause within-household and across-household correlation in claims and choice. To isolate the within-household correlation, equation (2) is re-estimated with household-fixed effects, and Table 4 columns (2) and (4) display the results. The positive and significant point estimate for α indicates that households have more claims when the households select more coverage. In addition to the sign, the magnitude of α is notable: the within-household coefficient is about 60% of the overall coefficient, meaning a substantial amount of the positive correlation in this environment is within-household correlation, over time.[24]

[19] Comparing the samples used in these two specifications, we see that some households do not select the same coverage tier for dental and medical insurance. These households (those with a discrepancy in chosen coverage tier at any point in the data) represent 16,185 of the 116,426 household-year observations.

[20] An additional assumption is needed to calculate $Claims_{h,t}$ under the first definition of $Choice_{h,t}$: for households that leave some potential dependents uninsured, it is assumed that uninsured dependents have no claims because spending data are not observed for uninsured dependents.

[21] $Claims_{h,t}$ is calculated by applying the Plan H cost-sharing rules to the spending of each household, taking into account the different types of spending the household did (beyond just the average coinsurance rate).

[22] Examples of coarser measures of claims include a claim indicator or claim count. Using a measure similar to the one used here, Chiappori et al. (2006) suggest that the following condition is a test for positive correlation that is robust to many permutations in the plan differences, utility framework, and competitive environment: $\int R_H(d)\mathrm{dF}_H(d|X_i) \geq \int R_H(d)\mathrm{dF}_L(d|X_i)$, where $R_H(d)$ is the insurance payout for someone with d dental expenditures enrolled in Plan H and F_j, $j \in \{H, L\}$, is the distribution of dental expenditures for those who choose Plan j. When this condition holds, the authors say there is evidence of "relevant" asymmetric information. The test outlined in equation (2) is simply this condition where the conditioning on X_i in the expectation is constrained to be linear.

[23] When the second definition of $Choice_{h,t}$ is used, dental coverage tier is also included as an exogenously priced household characteristic in $X_{h,t}$ since the sample in this specification is restricted to households that treat their family composition as fixed for the purpose of insurance enrollment. Since the first definition of $Choice_{h,t}$ includes both plan and coverage tier choice, in this specification only the premium menu is controlled for as the whole premium menu (including premiums for multiple coverage tiers) may influence this definition of choice.

[24] One might be worried that the difference between the within- and across-household estimation is due to composi-

These estimates reveal that households switching from Plan L to Plan H increase their claiming behavior after switching (or correspondingly, households decrease claiming behavior after switching from Plan H to Plan L). Though ex post adverse selection is one explanation for this association, this evidence alone does not allow one to assess how much of this within-household, dynamic asymmetric information is driven by ex post adverse selection as opposed to other forces. This motivates my second approach (outlined in Section 5) in which I develop a model that links the strategic delay of treatment to the subsequent incentive to select more insurance coverage. This model allows me to then investigate the relative importance of ex post adverse selection as compared to traditional adverse risk selection through counterfactual analysis.

4.2 Strategic Claim Delay

To look for evidence of strategic claim delay, the annual maximum benefit feature of the dental plans is exploited. When individuals are close to exhausting the annual maximum of their plan, they may have an incentive to delay treatment until the benefits reset next January. I inspect monthly dental spending in adjacent years for this predicted pattern of elevated beginning-of-year expenditures for individuals who are incentivized to delay claims. The sample here is limited to those individuals who were covered under the company dental benefits during adjacent years, who had positive expenditures across these years, and who were enrolled in Plan L during the first of these years. Because individuals with more expenditures across the two years are more likely to have been incentivized to delay claims between the years, spending patterns are separately examined for those with a lot of expenditures across the two years (expenditures exceeding \$1,400), and those with less expenditures across the two years (expenditures less than \$1,400). Figure 3 plots the average fraction of total individual expenditures by month, and this fraction is normalized to one so that a flat line at one would indicate that spending is, on average, equally distributed across the months. Among those more incentivized individuals, there is a large increase in the fraction of spending beginning in January of year 2, and this elevation in spending persists for the following six months. For these incentivized individuals, on average 32% of the total dental spending across the two years is incurred during the first six months of year 2. In addition, there is an associated relative dip in spending among these people at the end of year 1. The patterns in this figure suggest that incentivized individuals delay some treatments from the end of year 1 to the beginning of year 2 to get more coverage for these treatments. There is comparatively little monthly variation in dental expenditures among the less incentivized individuals. In Appendix A, I demonstrate that these patterns are qualitatively unchanged when the cutoff between low- and high-spending individuals is moved by a few hundred dollars or when controlling for year 2 dental coverage.[25]

The patterns in Figure 3 align with the incentives to delay claims in this context. When an in-

tional differences between the households that are in the sample for longer or short periods of time (perhaps because of differing job turnover rates). This is not the case. In Appendix A, I demonstrate that the results are qualitatively similar when the sample is restricted to those households that remain with the company and enrolled on a relevant plan throughout the data period.

[25] As one would expect, those individuals that switch from Plan L to Plan H between the adjacent years display higher displacement between the end of year 1 and the beginning of year 2 than the no-switchers display. Still, those who do not switch also display displacement though somewhat smaller in magnitude. Of course, even non-switchers have incentives to delay in this context because of the annual maximum benefit feature of coverage.

dividual with spending close to the maximum benefit learns of a problem that requires treatment that would exceed this maximum, he has an incentive to delay this treatment until his benefits reset next January. In this figure, we can clearly see that many individuals respond to incentives to delay treatment until just after the commencement of a new calendar year. To explicitly link this strategic delay of claims to the adverse selection that this delay can cause, I develop a model in Section 5 that allows for the investigation of counterfactual policies, such as contracting on pre-existing conditions and open-enrollment periods. This model accounts for rigidities that may prevent one from delaying treatments when incentivized. In this model, correlations similar to those displayed in Figure 3 identify these delay rigidities.

The annual maximum benefit encourages some individuals to delay claims until the benefit resets upon the commencement of a new calendar year. In addition, individuals may want to delay more treatment if they can switch to more generous insurance coverage in the following year. Because individuals may have an incentive to delay more claims when they can subsequently select more insurance coverage, the pattern of monthly spending should depend on whether insurance coverage is locked during the two adjacent years (the annual maximum resets to zero but a person cannot choose a more generous policy) or whether plan switching is permitted during the two years (the annual maximum resets and a person can choose a more generous policy). This pattern, if it exists, should be mostly concentrated among those individuals with very high dental spending over the two years.[26] Thus, I use $2,000 as a cutoff to define incentivized individuals in Figure 4. Panels (a) and (b) of this figure plot the series separately for the periods when insurance decisions are unlocked and locked, respectively. In both panels, we see a more dramatic pattern than in the previous figure as these high-spending individuals are more likely incentivized to delay claims. Comparing Panel (a) to Panel (b), the fraction of total spending over the first six months of year 2 is substantially larger in the unlocked period than in the locked period. In particular, the elevated spending during the unlocked period persists for the first six months of year 2 while in the locked period the corresponding elevation in spending is shorter in duration. In the unlocked periods as compared to the locked period, on average 7% more of the total spending across the two years is done in the first six months of year 2. Overall, Figure 4 shows that both the annual maximum benefit feature and the opportunity to switch coverage seem to encourage people to delay treatment in this setting.

5 Empirical Model: Setup, Identification, Estimation, and Results

The previous section reveals direct evidence from the data consistent with strategic treatment delay and dynamic asymmetric information. From this evidence alone, however, it is not possible to discern how much adverse selection is driven by timing incentives (as opposed to other sources of selection). Thus, I develop and estimate a model that precisely specifies the determinants of the decision to delay claims links this to subsequent incentive to select more insurance coverage. This

[26] These high-spending individuals are more likely to have received problems in year 1 for which the treatment costs would have far exceed the year 1 maximum benefit. Thus, for these high-spending individuals, delaying claims is potentially much more appealing when they are able to switch to Plan H to face a larger maximum benefit.

model makes explicit the costs of delay (which may include pain and cognitive costs). Using this framework, I estimate unobserved heterogeneity in dental risk and delay costs. Section 6 then uses these estimates as inputs to investigate the relative importance of ex post adverse selection and to explore counterfactual policies.

5.1 Model Setup

5.1.1 Overview

Though ex post adverse selection is of central importance in the counterfactual analysis, the focus of the model and estimation is on the economic primitives that affect the decision to delay claims between two adjacent years. This modeling approach takes advantage of two characteristics specific to this environment. First, the two-year period during which the company locked insurance coverage allows me to estimate the model without extra assumptions or data requirements that would be needed to model endogenous plan switching. Using data from the locked period, 2005-2006, the model is used to estimate unobserved heterogeneity in delay costs and dental risk abstracting from plan choice and plan switching.[27] Second, the incentive to switch plans in this environment is aligned with the incentive to delay claims from one year to the next. Though the model is estimated using the locked period during which there is no insurance plan switching, this second feature means that the estimated delay frictions are able to capture some of the frictions that may inhibit optimal insurance plan switching more generally. Under some additional assumptions, which are made clear later in the paper, the estimated distributions of dental risk and delay costs are sufficient for policy analysis related to claim timing and ex post adverse selection.

I first outline the model for a single individual, and then explain how heterogeneity is incorporated. In the model, the individual decides how much treatment to postpone after realizing events (in this case, dental problems). To be clear on the terminology used here, the "timing of events" or the "event date" will refer to the date the individual becomes aware of an event and the necessary associated treatment.[28] The timing of the model is outlined below.

1. The individual realizes first-year dental events.
2. The individual decides how much treatment in dollars to delay until the second year, m. He makes this decision according to his cost of delaying claims, $c(\alpha_i, m)$. He treats the remaining events for which he has not delayed treatment.
3. In year 2, the individual realizes new events, and he treats these new events along with any treatment delayed from the previous year.

In this baseline model, claims (treatments) cannot be delayed beyond year 2, and claims are not carried over to year 1 from prior years. In other words, I abstract from initial and terminal condition

[27] Because plan switching decisions rely on individuals' expectations about plan offerings many years into the future, modeling such behavior with a short panel would require fairly heroic assumptions about expectations into the distant future. For this reason, I focus the delay decision between adjacent years during which insurance coverage is locked. In contrast to the plan switching decision, this delay decision requires very little foresight and thus it is possible to convincing capture all the relevant information in a stylized model using the short panel that is available.

[28] It is the date of event recognition that is important for strategic delay, not the date the problem was initially acquired (if different from the date of recognition).

issues.[29] While this simplifying assumption is employed in the baseline specification, Appendix C describes and estimates an alternative specification that relaxes this assumption by allowing individuals to consider the possibility of delaying claims in the future when deciding on the amount of claims to delay in the present. This alternative specification yields results similar to the baseline estimates.

There are two components of the empirical model: the cost model and the decision model. The cost model specifies the frequency of events and costliness of treatment, while the decision model specifies the incentives to postpone treatment from year 1 to year 2.

5.1.2 Cost Model

The cost model describes both the frequency of dental events and the financial cost associated with treating these events. When individuals strategically delay treatment, the timing of events diverges from the timing of claims. Because the identification of the model relies on quantifying the extent of this divergence, it is important to describe the underlying frequency of events. One major challenge in this setting is that event timing is not directly observed as the data contains only claim information available to the insurer. Though the timing of claims puts some bounds on the timing of events, I need to place some additional restrictions on the timing of events in order to estimate delay costs.

The assumption I make is that dental events arrive independently over time where the rate of arrival depends on the individual's "risk type," λ_i.[30] This conditional independence assumption implies that the number of events received by the individual in year t, n_t, is governed by a Poisson distribution[31]: $n_t \sim \text{Poiss}(\lambda_i)$.

The individual's decision to postpone claims will depend on the number of events he expects to receive the following year and the expense of treating these events. For each event l, the treatment cost, c_l, is assumed to be an independent draw from the "cost intensity" distribution, G, representing the empirical distribution of claim costs in the data, $c_l \sim$ G.

In the estimation, this distribution is allowed to vary with age in a categorical manner to capture the fact that the types of treatments done by the middle-aged differ from those done by children.[32]

[29] Also, implicit in this setup is the assumption that the only form of moral hazard is delaying treatment. That is, there are no optional treatments (or treatments that may be delayed indefinitely) in the model. In reality, some dental procedures are very optional. This assumption is made to focus on claim timing rather than purely static moral hazard. (The most obvious example of optional dental procedures are those done for purely cosmetic reasons that have become increasingly popular in recent years. Because cosmetic dentistry is not covered by the dental insurance offered by the company, these procedures are excluded from the analysis.)While this assumption may exclude some forms of moral hazard, much of what is traditionally thought of as moral hazard in this context has a timing element as well. Intuitively, if an individual does an "extra" procedure today, this crowds out the probability that he will do the same procedure for the same problem tomorrow. Returning to the two-period model, the assumption of no optional treatment implies that treating events is mandatory. An individual will treat his events when it is advantageous to do so, but ultimately he will treat them within the two years.

[30] It is important again to highlight the distinction between events and claims. Although events (e.g., cavities) are assumed to be independent over time, claims (e.g., fillings) can be serially correlated for a number of reasons including individuals' decisions to delay claims based on insurance incentives.

[31] Though at first blush this independence assumption may seem restrictive, the model is robust to some correlation in events. The model accommodates cross-sectional correlation in events through allowing for rich heterogeneity in risk types.

[32] For this purpose, age is discretized into the following categories: Less than 19 years of age, 19-30 years of age,

Using this notation, the total cost of treating the events that arrive for the individual in year t can be written as follows:

$$d_t = \Sigma_{l \leq n_t} c_l. \tag{3}$$

5.1.3 Decision Model

The decision model describes how the individual chooses the amount of treatment to delay between the adjacent years. Suppose the individual is enrolled in Plan j over the adjacent years, and the individual has Constant Absolute Risk Aversion per-period utility, $u(c)$, with risk aversion r.[33]

The individual maximizes the sum of his first-year utility and his expected second-year utility by choosing m, the dollar value of delayed claims, from M, the set of possible delay amounts. The individual does not discount his second period utility, and the individual is permitted to save and borrow without interest, denoted as s below.[34] The individual's optimization problem is below:

$$\max_{m,s} \quad u(-OOP_j(d_1 - m) - c(\alpha_i, m) - s) + E_{\tilde{d}_2}[u(-OOP_j(\tilde{d}_2 + m) + s)|\lambda_i, G] \tag{4}$$

where $m \in M$.

In the first year, the individual pays out-of-pocket expenses associated with his incurred dental claims, $\mathrm{OOP}_j(d_1 - m)$. The individual also pays a "delay cost" to postpone m dollars worth of claims, $c(\alpha_i, m)$, and the individual is permitted to save, s. In the second year, the individual pays dental out-of-pocket expenses that are a function of the delayed claims and the second year events, $\mathrm{OOP}_j(d_2 + m)$. In addition, he receives savings from the previous year, s. When the individual makes his decision to postpone treatment, he does not yet know what dental events will arrive in the second year or how costly these events will be to treat. Thus, the second term in equation (4) is the expected second-year utility where the expectation is taken over d_2, the cost of treating events that arrive in year 2. This expectation is conditional on the information the individual knows about the distribution of d_2: his risk type, λ_i, and the cost intensity distribution, G.

The individual may find it costly to delay treating dental events for many reasons, including physical pain caused by delaying a procedure, discomfort with postponing a procedure recommended by a dental professional, and inconvenience associated with maximizing insurance coverage of dental needs. To rationalize these frictions, I assume the individual pays the following delay cost:

$$c(\alpha_i, m) = \alpha_i \mathrm{I}(m > 0). \tag{5}$$

This cost is equal to α_i for any positive amount of delayed treatment, and more structure is added to this delay cost when incorporating heterogeneity.

31-40 years of age, 41-50 years of age, and over 50 years of age. These are the same categories across which the average coinsurance rate varies (used in the out-of-pocket cost function $\mathrm{OOP}_j(x)$).

[33]The CARA per-year utility form is convenient as any component of consumption that remains constant across the two years will be ignored in the individual's optimization problem. I assume the individual does not expect changes in income over the two years so income drops out of the individual's decision to postpone dental claims. Although premiums change slightly over the locked period ($5 on average), I assume that individuals treat premiums as constant across the two years as well.

[34]Savings is assumed to be bounded by the wage net current period expenses ($\bar{s} = w - \mathrm{OOP}_j(d_1 - m) - c(m, \alpha_i) - p_1$) and borrowing is restricted to be smaller than the individual's wage minus the out-of-pocket expenses in the worst cost draw in the population in year 1 ($\underline{s} = p_2 + \mathrm{OOP}_j(\bar{d}_1) - w_i$). These constraints are non-binding in practice for the individuals in the sample.

If individuals could delay any continuous amount of claims, this model would predict a large mass of annual dental spending corresponding to the level that exactly exhausts the maximum benefit. In the data, however, many individuals have dental spending close to exhausting the maximum benefit, or dental spending just above the maximum benefit. This pattern suggests that individuals face some rigidities in the dollar value of treatment they may delay, so it is conceptually important for the model to include some restrictions on the possible delay amounts, M, and I make two such restrictions.[35] It should be emphasized that although it is conceptually important to make some restriction, the particular restrictions I make may not correspond to the actual restrictions individuals face; I make these particular restrictions only for computational simplicity.[36] First, I assume the individual can delay expenses only claim by claim. For example, this prevents the individual from delaying half of a cavity filling from year 1 to year 2. Second, I assume the individual cannot manipulate the claim order. This assumption means that the sequence of claims in the data reflects the sequence of events the individual received.

5.1.4 Heterogeneity

Unobserved heterogeneity is introduced in two parameters: the risk type, λ, and the delay cost α. The cost intensity distribution, on the other hand, varies with age in a categorical manner. This means the ex ante heterogeneity in total dental spending (conditional on age) stems from heterogeneity in the frequency of events (which is governed by the risk type λ), and not in the cost per event. In this context, the risk type can be thought of as a measure of one's overall dental health (conditional on age), where those with better health (lower risk) receive fewer events on average.

The parameters λ_i and α_i are known to the individual but unobserved by the researcher. The delay cost, α_i, is assumed independent of the risk type, λ_i. Risk type, λ_i, is assumed to come from a lognormal distribution with parameters $\mu_\lambda(X_i)$ and σ_λ:

$$\lambda_i \sim \text{lognormal}(\mu_\lambda(X_i), \sigma_\lambda). \tag{6}$$

The distribution of risk types is allowed to vary with individual characteristics, X_i, including age and gender,

$$\mu_\lambda(X_i) = X_i\beta. \tag{7}$$

For the delay cost, a more restrictive form of heterogeneity is assumed. In particular, I assume that individuals can either freely delay claims ($\alpha = 0$) or cannot delay claims ($\alpha = \infty$) between the two years. Thus, the delay cost distribution can be summarized by the fraction, p_α, of individuals who can freely delay claims:

[35] Absent any restrictions on M, a mass of people exactly at the benefit exhaustion point would be predicted because of the form of the cost function above. That is, because the cost function is not increasing in the amount of treatment delayed, individuals will either delay everything beyond the maximum or nothing. Practically, it is possible to assume a different cost functional form instead of adding additional rigidities in the possible delay amounts, M, to rationalize the data. However, restricting M in addition to the cost function captures the separate forces that lead to the degree of bunching we see in the data: the degree to which people can delay treatment for intrinsic reasons related to personal or procedure characteristics (captured in the cost function) and the limitations on bunching due to the fact that procedures are not divisible (captured by the restrictions on M).

[36] These particular restrictions simplify estimation through reducing the size of the individual's choice set.

$$\alpha_i = \begin{cases} 0 & \text{with probability } p_\alpha \\ \infty & \text{with probability } 1 - p_\alpha. \end{cases} \tag{8}$$

There are multiple ways to interpret the estimated heterogeneity in the delay cost. The variation in delay costs could be driven by individual characteristics such as sophistication in navigating insurance incentives or pain tolerance. Alternatively, the variation in delay costs could be driven by characteristics of marginal procedures such as urgency. In reality, the cost of delaying claims probably has some component that is procedure-specific but common across individuals and some component that is individual-specific but common across procedures. The two-period model cannot separately identify the individual-specific and procedure-specific sources of heterogeneity in delay costs. Instead, the estimate of delay cost heterogeneity can be interpreted as a combination of both sources of heterogeneity. Counterfactual policies are evaluated in Section 6 under two extreme interpretations of this delay cost heterogeneity.

5.2 Identification

The risk type distribution parameters, $(\beta_\lambda, \sigma_\lambda)$, are identified by the variation across individuals in the total number of claims observed during the two years (without regard for year-to-year timing). The delay cost distribution parameter, (p_α), is identified by the division of claims between the two years among individuals with claims close to the maximum benefit in the first year (who likely had an incentive to delay claims). Notice this source of identification is very similar to the evidence depicted in Figure 3 (discussed in Section 4).

Intuitively, only individuals with the incentive to delay claims give us any information about the delay costs. In this environment, individuals have an incentive to delay claims if and only if they receive year 1 events that put them beyond the maximum benefit. Because events are not observed directly, I must use the model, in combination with the claims data, to infer which individuals may have received events that would have put them at the plan maximum benefit in the first year. These "potential delayers" are those individuals who have first-year claims that either exceed the maximum benefit or would have exceeded the maximum benefit if their first claim from year 2 had instead been claimed in year 1.[37] In the sample used to estimate the model, 2.2% are potential delayers.[38] For those who are potential delayers, the model determines the likelihood of observing the data given the possible delay costs.[39] Intuitively, an increase in the fraction of potential delayers with many

[37]The identification of these "potential delayers" relies on the definition of possible delay amounts (M). Because the definition of M implies the sequence of claims is fixed, this allows one to conclude that an individual did not consider delaying claims if the sum of his first-year dental spending plus the first claim in year 2 does not reach the maximum benefit. To illustrate this point, I suppose the opposite and show this leads us to a contradiction. In particular, suppose that such an individual did delay claims. Then, in the first year this individual must have received the event associated with the first claim in year 2, since the sequence of claims is fixed. However, since this claim would not have pushed him over the maximum benefit, it would have been suboptimal to delay this claim to year 2 and give up benefits, leading to a contradiction as this is inconsistent with the data. Thus, according to the model, only "potential delayers," as defined in the text, may have considered delaying claims.

[38]Among individuals over 18 years of age, 3.1% are potential delayers.

[39]The likelihood of observing an individual's claims given that he can freely postpone claims is simply the probability that the individual received any combination of events that would have led to the observed division of claims between the two years given optimal behavior and costless delay. On the other hand, the likelihood of observing an individual's claims given that his delay cost is infinite is simply the probability that his division of events between the two years mimics his division of claims observed in the data. The delay parameter is estimated by maximizing the likelihood of

claims beyond the year 1 maximum benefit would lead to a lower estimate of p_α (the fraction that can freely delay claims if incentivized).

In the counterfactual analysis, the estimated delay probability is applied to the entire sample. Though the delay parameter is identified by these incentivized individuals, the estimated delay cost distribution describes the parameter relevant in the entire sample under two assumptions: (1) independence between risk types and delay costs and (2) plan enrollment is unrelated to delay costs. In Appendix C, I show the model estimates and the counterfactual analysis are robust to relaxing this second assumption. Although it is empirically infeasible to relax the first assumption, any positive correlation between the delay cost and risk type would lead the model to underestimate the propensity to delay claims.[40]

5.3 Estimation

Two sets of parameters are estimated: the distribution of risk types, parameterized by $(\beta_\lambda, \sigma_\lambda)$, and the distribution delay costs, parameterized by p_α. These parameters are estimated taking the empirical cost intensity distribution, G, as given.

Let $\Theta = (\beta_\lambda, \sigma_\lambda, p_\alpha)$. Define $claims_i$ as the sequence of claims observed for individual i over the two years, and define D_i as the division of this sequence of claims into those done in year 1 and those done in year 2. The individual's contribution to the likelihood of Θ can be written as follows:

$$l_i(\Theta|claims_i, D_i) = \int P(claims_i, D_i|\lambda, \alpha) dF(\lambda, \alpha|\Theta). \tag{9}$$

Because λ_i and α_i are known to the individual but not observed by the researcher, the likelihood must integrate over the distribution of these latent parameters, $F(\lambda, \alpha|\Theta)$.

Aggregating across individuals, the complete likelihood for Θ can be written as follows:

$$L(\Theta|\text{claims,D}) = \prod_i l_i(\Theta|claims_i, D_i). \tag{10}$$

The method of Maximum Simulated Likelihood is used to estimate Θ. The details of this estimation are outlined in Appendix B.

The coefficient of absolute risk aversion enters this calculation through the expected utility calculation associated with delaying treatment. Risk aversion is assumed to be common across individuals.

the parametric distribution given the observed division of claims among these potential delayers.

[40] In the estimation, the delay cost, α_i, is assumed independent of the risk type, λ_i. Though ideally the empirical model would allow for correlation in these parameters, in practice it would be difficult to identify this correlation. The reasoning behind this is quite simple. The identification of the delay cost parameter comes from those individuals who are on the margin of delaying claims from one year to the next because they are close to the maximum benefit in the first year. Of course, those who are close to the maximum benefit are relatively risky (have relatively high risk types, λ_i). It is difficult to uncover the propensity to delay claims among low risk individuals as they are typically far from the maximum benefit giving them no incentive to delay claims. Thus, it is not possible empirically identify how the delay probability, p_α, varies over a large range of risk types. Though it is not possible to empirically identify the correlation between the delay costs and the risk types, any correlation that exists in reality would affect the interpretation the estimated delay probability, p_α. Consider the case when the delay cost is positively correlated with the risk type ($\text{Corr}(\alpha_i, \lambda_i) > 0$), meaning that those with high risk types are less able to delay treatments than those with low risk types. One way to generate this sort of correlation would be if those with high risk types had a higher propensity to suffer from urgent dental problems. If this is the case, the estimated delay probability will underestimate the average delay probability in the population. On the other hand if the delay cost is negatively correlated with the risk type ($\text{Corr}(\alpha_i, \lambda_i) < 0$), the estimated delay probability will overestimate the average delay probability in the population.

The risk aversion value used in the baseline estimation, 2.3×10^{-3}, is calibrated to match the observed plan shares in the data had the households made a static insurance decision.[41] Contrary to reality, if households could not delay claims and could not self-insure through savings, the calibrated risk aversion value would imply simulated plan shares that are roughly equivalent to the plan shares observed in the data.[42] However, because the ability to delay claims and the ability to self-insure through savings/borrowing both make insurance less attractive in a given year, the calibrated value of risk aversion underestimates the true risk aversion necessary to justify the plan choices in the data. In Appendix C, I estimate the model under different risk aversion values and show the basic lessons of the counterfactual analysis are qualitatively unchanged.

5.4 Results

The parameter estimates from the model are displayed in Table 5, along with bootstrapped standard errors. The parameters $(\beta_\lambda, \sigma_\lambda)$ describe the heterogeneity of risk types in the sample. Appendix A illustrates that the implied annual spending fits the data quite well. The heterogeneity in delay cost is described by p_α, the fraction of individuals in the population who can freely delay claims from one year to the next. The data suggest that a large fraction of people respond to insurance incentives by delaying claims. Approximately 48% of individuals may freely delay claims when incentivized, according to the parameter estimates. Because the estimates suggest individuals can plan and delay dental treatment quite easily, the associated ex post adverse selection has the potential to be quite severe.

Though individuals cannot switch insurance coverage during the two-year coverage period used to estimate the model, in general individuals who have postponed claims may have an incentive to switch to better insurance coverage ex post. Of course, frictions may prevent this optimal ex post switching as previous studies have found substantial switching costs in related contexts.[43] In the context of claim timing, frictions that prevent optimal plan switching can lessen ex post adverse selection and improve welfare in equilibrium. Because the model focuses on the delay decision in isolation, I do not specify or estimate frictions that prevent employees from selecting plans wisely or switching plans when appropriate. However, the estimated degree of claim delay may reveal some information about switching frictions in this setting. Some typical explanations of suboptimal plan switching are that individuals lack information about plan details or lack the cognitive sophistication necessary to navigate complex incentives. These explanations and many others used to explain switching costs are also potentially important explanations for the suboptimal delay of claims. A feature of this context is that the only motivation to buy more coverage comes from delaying treatment, so long as premiums remain stable and risk types are fixed. In the data used to

[41] Following Cohen and Einav (2007), one can interpret the calibrated risk aversion value as follows: an individual with this risk aversion would be indifferent between taking a gamble of winning \$100 and losing \$81 each with probability one half and not taking this gamble.

[42] In this calibration, it is assumed that premiums are tax-free and the marginal tax rate is 25% for all in the sample. If premiums are treated as post-tax expenditures, it would take an even higher value of risk aversion to rationalize the plan shares.

[43] Many studies have found substantial switching costs in related contexts (e.g., Handel 2010, Madrian & Shea 2001). One might expect that incentives to switch insurance plans subsequent to delaying treatment are more salient to individuals than simple changes in prices or defaults, for example. Thus, we may expect individuals to face limited switching frictions conditional on delaying treatment in response to insurance incentives.

estimate the model, this means that individuals who had an incentive to buy better coverage after delaying treatment (though could not because coverage decisions were locked), also had an incentive to delay treatment even though insurance coverage was fixed. If the frictions that prevent one from optimally switching plans also prevent one from optimally delaying treatment, then the estimated fraction of individuals who optimally delay claims is a lower bound on the fraction of individuals who, given the choice, would have switched plans if it were optimal to do so.[44] In the following section, I make some assumptions about the relationship between delay costs and switching costs to analyze counterfactual policies.

6 Counterfactual Analysis

The parameter estimates are used to investigate the impact of strategic timing and the associated ex post adverse selection on insurance enrollment and insurer costs. In the interest of exploring broader questions relating to the overall insurance market, the counterfactual analysis focuses on the impact of adverse selection outside the scope of the insurance options available within the firm. In particular, the counterfactual analysis investigates the market for comprehensive dental insurance, a product that in practice does not exist, to explore potential explanations for this market's unraveling. I compare insurance enrollment and insurer costs in four different contracting scenarios: insurers price no information (as is typical in employer-provided insurance), insurers price risk types, insurers price pre-existing events, and insurers price both pre-existing events and risk types. This analysis allows me to isolate the impact of ex post adverse selection (resulting from the non-contractibility of pre-existing events), and compare the impact of this selection to the impact of traditional adverse risk selection (resulting from the non-contractibility of ex ante risk types). Overall, I find that either ex post adverse selection or traditional adverse risk selection is severe enough in this setting to cause the market for full coverage insurance to largely unravel. In addition, I investigate the impact of enrollment frequency restrictions on the viability of comprehensive dental insurance by comparing annual enrollment to less frequent enrollment. I find that restrictions on the frequency of open-enrollment periods, when coupled with modest premium subsidies and risk-rating, may substantially increase insurance enrollment.

Before continuing, it is important to highlight that the model abstracts from any potential benefits from choice as modeled individuals have homogeneous preferences (in this case, risk aversion), and the model does not include insurer competition. Thus, the focus of the counterfactual analysis is to shed light on the costs of choice relative to the first best of full insurance. It is also important to keep in mind that the model estimates come from a particular insurance setting and a non-representative population; thus, we should exercise the appropriate amount of caution in interpreting the counterfactual analysis based on these estimates.

[44] The logic behind this point is simple. Suppose that the only frictions that prevent one from optimally switching plans (for example, the lack of cognitive ability necessary to understand insurance incentives) also prevent one from optimally delaying treatment. Then, individuals who are sophisticated enough to delay claims are also sophisticated enough to switch insurance coverage if it is optimal to do so. Some other individuals (who could not delay claims when incentivized because of pain or other reasons) may also be sophisticated enough to switch coverage if it were optimal. Thus, in this case, the fraction of individuals who can delay claims optimally is a lower bound on the fraction of individuals who can optimally re-evaluate insurance coverage.

6.1 Setup

Below, I explain how the model is modified to look at the broader counterfactuals of interest. The time horizon of the counterfactuals is extended to ten years.[45] In each period, individuals may select whether to buy full coverage dental insurance or go without insurance.[46] In each counterfactual scenario, I simulate dental events, timing decisions, saving decisions, and coverage decisions, assuming individuals make these decisions by maximizing the sum of their current year utility and their future expected utility. Thus, individuals solve for the optimal insurance coverage in each period, taking into account the option to delay claims, save/borrow, and select insurance coverage in future periods. Constraining premiums to be constant across the ten years of available insurance, I solve for the equilibrium numerically both when insurers break even and under various premium subsidies.[47]

A key parameter in the model is p_α, the fraction of individuals who freely delay claims between adjacent years.[48] In the counterfactuals, I consider two extreme cases for this delay cost heterogeneity. In the first case, I assume the delay cost is "individual-specific," meaning each individual's delay cost, α, is fixed across time and the individuals know their delay cost ex ante. One possible interpretation of this case is that some individuals are sophisticated while others are not, but the procedures are all equally urgent. In this case, I assume that individual characteristics are the only impediments to claim delay and that these same individual characteristics are the only impediments to optimal plan switching. In other words, the fraction of the population that optimally delay claims also optimally switch insurance plans. In the second case, I assume all the heterogeneity in delay is "procedure-specific" in that individuals receive an independent draw from the delay cost distribution in each year along with their draw of events (after their insurance coverage decision in that year). In this case, it is unknown to an individual ex ante whether he will be able to delay claims in a given period. The interpretation of this case is that individuals have no inherent characteristics that make them more or less prone to delay claims, and in this case, I assume that there are no frictions for optimal plan switching.

Aside from the delay cost, there are no restrictions on an individual's ability to delay claims for multiple years. In particular, in all the counterfactual simulations, individuals can delay claims for multiple years so long as the delay cost is zero in each relevant year. In the case of individual-specific delay costs, this means that individuals with zero delay costs can continue to delay claims until the last year of the simulation. In the procedure-specific case, this means that individuals can

[45]This is a departure from the estimation of the empirical model where the relevant time horizon was only two years.

[46]For simplicity, the counterfactual analysis focuses on individual-level decisions as opposed to household-level decisions. Modeling individual-level decisions allows one to avoid specifying within-household correlation structures on risk types and delay propensity.

[47]Because individuals choose insurance conditional on future insurance premiums, the premium for full coverage insurance is assumed to be fixed over time and exogenous with respect to individual decisions.

[48]In the counterfactuals, two options are available to individuals: no insurance or full insurance. With these options, the incentive to delay claims comes solely from the opportunity to buy full coverage insurance after delaying treatment (as opposed to the insurance nonlinearities used to estimate the model). This means that even individuals with low dental expenditures may decide to delay claims if they are uninsured in the current period. Though the delay cost parameter is estimated using variation in insurance incentives near the annual maximum benefit of the company dental plans, in the counterfactual analysis this delay parameter is used to evaluate optimal delay decisions in all parts of the dental spending distribution.

potentially delay treating events for multiple years, but there is a 52% chance in any given year that the individuals' accumulation of events will become urgent and will need to be treated immediately.[49]

6.2 Information Contractibility

The first set of counterfactuals explores the market for comprehensive dental insurance in different contracting settings. Individuals are assumed to have two insurance coverage options in each year: full insurance or no insurance. I simulate annual coverage decisions and claims under four different contracting scenarios: the insurer contracts on no information (the typical situation in employer-provided insurance), the insurer contracts on risk type but not pre-existing events, the insurer contracts on pre-existing events but not risk type, and the insurer contracts on both risk type and pre-existing events. When the insurer contracts on risk types, the insurer and the individual have symmetric information about the individual's ex ante risk type, λ, and the insurer can price policies conditioning on this information. If pre-existing events are contractible, the insurer can observe and price events pre-existing at the time of insurance enrollment; in this case, individuals have no incentive to delay claims.[50] I numerically compute an equilibrium when the insurer breaks even, subsidizes premiums at 25%, subsidizes premiums at 50%, and subsidizes premiums at 75%. Details of this computation are given in Appendix B.

Table 6 displays the percent of insured individual-years, as well as the average insured cost in each of the considered scenarios. One can see that the unsubsidized market for full insurance suffers from severe adverse selection unless both risk types and pre-existing events are contractible. Insurance enrollment is lowest when insurers contract on neither ex ante risk types nor pre-existing events (as is common in practice). When there is no premium subsidy or when there is a modest premium subsidy, it appears as if ex post adverse selection (due to the non-contractibility of pre-existing events) depresses coverage more than traditional adverse selection (due to the non-contractibility of risk types). For example, when pre-existing events are not contractible but risk types are, 0.6% and 2.0% insure when delay costs are procedure-specific and individual-specific, respectively; in either of these cases, 9.1% would insure if instead events were contractible but risk types were not. However, this pattern does not hold under larger subsidies and procedure-specific delay costs, in which case traditional adverse risk selection depresses insurance coverage more.

In general, adverse selection appears to be more severe when delay costs are individual-specific than when delay costs are procedure-specific. In other words, strategic timing is a larger problem in this market when delay costs are persistent over time as opposed to randomly reassigned each year. The economic rationale for this is simple. When the ability to delay claims is tied to persistent individual characteristics, sophisticated individuals have a lot of freedom to plan and time treatments and to go without insurance in most periods, knowing they have this flexibility. In contrast, when it is uncertain whether an individual will be able to postpone treatment, insurance is more attractive because of the possibility of urgent procedures. The findings in Table 6 show that subsidies are

[49] In this case, there is a $(1 - p_\alpha) = 52\%$ chance that individuals will need to treat their accumulation of events in a given year.

[50] The ability to condition premiums on pre-existing events can be conceptually thought of as the situation in which the insurer can back-date individuals' dental problems and condition premiums on this information.

more effective at promoting insurance coverage when delay costs are less persistent. In addition, the relative effectiveness of subsidies in promoting insurance enrollment depends on the persistence of delay costs. When delay costs are randomly reassigned over time (procedure-specific), it seems that large subsidies are more effective at promoting coverage in the presence of ex post adverse selection than traditional adverse risk selection. In contrast, when delay costs are very persistent over time (individual-specific), subsidies are more effective at promoting coverage in the presence of traditional adverse risk selection than ex post adverse selection.

Overall, if one extrapolates from these findings, it seems that ex post adverse selection is severe enough in this setting to explain the lack of annual comprehensive dental insurance seen in practice. In other words, this ex post adverse selection may explain why so few people have dental policies in the US and why the available policies offer no coverage for the right tail of dental risk. Because dental risks are typically not "catastrophic" and individuals may self-insure to some degree by saving/borrowing, insurance may be less important in this context than in other contexts. Still, the value of comprehensive insurance is non-trivial in this setting. Using the parameter estimates, I find the total surplus that would be generated by universal full insurance for dental risk (the first best in this model) relative to no insurance is at least 10% of total dental spending, or $10 billion nationwide annually.[51] Thus, the observed underinsurance of dental risk is likely associated with considerable welfare losses relative to the first best.[52]

6.3 Enrollment Period Frequency Restrictions

Restricting the frequency of insurance enrollment may limit adverse selection induced by strategic timing. I investigate four scenarios with different open-enrollment period frequencies: annual enrollment, enrollment every two years, enrollment every five years and lifetime enrollment. In the case of annual enrollment, for example, an individual decides on a contract just before the start of each year. In the case of enrollment every two years, an individual chooses a contract for the first two years just before the start of year 1, and so on. In each scenario, individuals have two options: to buy full insurance or no insurance. To isolate the impact of enrollment frequency on ex post adverse selection, I calculate an equilibrium in each scenario assuming the insurer contracts on risk types but not on pre-existing events. An equilibrium is calculated numerically when the insurer breaks even and when the insurer subsidizes premiums at 25%, 50%, and 75%. The details of this calculation are outlined in Appendix B.

Table 7 displays the fraction of individual-years insured and the average insured cost in each of the investigated scenarios. Regardless of enrollment frequency restrictions, the unsubsidized market

[51] Details of this calculation are in Appendix B. For reasons discussed in Appendix B, this is an underestimate of the value of full insurance in this setting.

[52] This estimated benefit of full insurance verses no insurance does not include any administrative costs. If administrative costs are non-negligible, it is not clear that the first best would be full insurance. In the counterfactual analysis, I abstract from administrative costs to make the point that ex post adverse selection is severe enough to explain the unraveling of a market for full coverage dental insurance even under the most optimistic of cases in which administrative costs are negligible. Of course, this analysis does not preclude the possibility of other barriers obstructing a market for comprehensive dental insurance (such as substantial administrative costs). If there are other barriers, the counterfactual analysis reveals that even if one were to eliminate other barriers, still this market would unravel due to ex post adverse selection.

suffers from severe adverse selection except in the extreme case of lifetime enrollment. In all but the lifetime enrollment case, less than 2.0% of individual-years would be insured when premiums are unsubsidized. Over some ranges of premiums subsides, restricting the frequency of enrollment periods has a non-monotonic effect on insurance enrollment in the unsubsidized market. That is, in some cases reducing the frequency of enrollment depresses insurance enrollment at moderate frequencies. Intuitively, it is possible for adverse selection to become more severe under less frequent enrollment because individuals must commit to insurance contracts for longer so those strategically delaying claims have typically delayed more claims. Of course under lifetime contracts, the first best is achieved as this eliminates the incentive to delay treatment and is equivalent to the symmetric information case.

Still, it seems that reducing the frequency of enrollment can encourage comprehensive coverage when contract periods are extended to five years. For example, when delay costs are procedure-specific and premiums are subsidized at 25%, restricting enrollment frequency from annual enrollment to enrollment once every five years increases insurance coverage by 72.5 percentage points. In summary, it seems that restricting choice frequency may boost insurance coverage when coupled with subsidies and risk-rating, and extending the length of open-enrollment periods would be more effective if delay costs are not too persistent over time (closer to procedure-specific than than individual-specific).

6.4 Robustness

To ensure that the estimates and counterfactual analysis are not too sensitive to the assumptions used to get the baseline estimates, the model is re-estimated under alternative assumptions on the relevant time horizon, risk aversion value, and plan selection. The details of these alternative specifications as well as the results from them are displayed in Appendix C. The estimated delay probability ranges from 0.37 to 0.56 across the alternative specifications. Using these alternative delay probability estimates, the information contractibility counterfactual analysis is repeated and the results are qualitatively unchanged (results displayed in Appendix C). The robustness of this analysis indicates that a market for comprehensive dental insurance would not be viable regardless of whether 37% or 56% of treatments are delayable.

7 Conclusions and Future Work

The strategic timing of claims can cause inefficiencies in insurance markets. Using claim-level data, I find clear patterns that suggest individuals strategically delay dental treatment when insurance incentives encourage them to do so. I then develop and estimate a model that explicitly links this strategic delay of claims to the adverse selection it creates. Through counterfactual analysis, I find that this strategic delay of treatment and the associated ex post adverse selection is one potential explanation for the missing market for comprehensive dental insurance. In addition, I show that ex post adverse selection, which stems from the non-contractibility of pre-existing events, can cause severe underinsurance even when insurers and individuals have symmetric information about individuals' ex ante risk. The counterfactual analysis also reveals that market unraveling is typically more severe when the propensity to delay is linked to persistent individual characteristics. Overall, I find that comprehensive insurance for dental risk is unviable without substantial reforms relative to typical insurance markets such as: pricing more risk information, significantly restricting insurance choice frequency (thereby increasing coverage periods), and expanding premium subsidies. More generally, my results indicate that severe adverse selection may arise in settings where insured costs are elastic with respect to timing and the timing of the underling risk is not contractible.

There are a number of ways to extend and improve the analysis in this paper. Specifically, there are at least two ways to enrich the model that would allow one to gain additional insights related to this source of asymmetric information. First, the empirical model model focuses on the timing of spending abstracting from traditional moral hazard in the level of utilization. An interesting extension would be to adapt the model to include traditional moral hazard (optional spending) in addition to this timing moral hazard.[53] A model that encompasses both notions of moral hazard could be used to empirically estimate the separate forces that contribute to the elasticity of health care spending with respect to insurance coverage. Second, the empirical model abstracts from the consequences of delaying care on the evolution of risk. It is likely that delaying treatment can cause an individual's overall health to decline and thereby increase his risk of receiving bad events in the future.[54] Thus, one natural way to extend the model would be to allow risk types to dynamically evolve when treatment is postponed; an extension in this direction could allow one to quantify the long-term health costs of delaying care.

The asymmetric information studied in this paper arises because of the combination of the non-contractibility of pre-existing events and the ability of individuals to postpone treatment (once they know it will occur). There are many natural questions related to this topic that can be explored in future work. For instance, the non-contractibility of pre-existing events can interact with per-period nonlinearities commonly seen in insurance contracts. One such feature, the annual maximum benefit of dental policies, allows me to identify the strategic timing of spending in this paper. Contract per-period nonlinearities that encourage strategic timing are seen in many insurance settings, often in the form of annual deductibles and annual out-of-pocket maximums. The optimal design of such contract features in the presence of asymmetric information when the timing of risk is not contractible is an important area for future research.[55]

Another promising topic for future work involves the bundling of treatments for the purpose of insurance coverage. While the main purpose of health insurance is to allow individuals to pool risk, much of this risk pooling can break down when individuals can re-evaluate insurance decisions frequently.

[53] Increasing the out-of-pocket price of care on the margin can lead individuals to either delay some care to a future date in anticipation of lower out-of-pocket prices (timing moral hazard) or forgo some optional care completely (traditional static moral hazard).

[54] The long-term health consequences of delaying care may motivate the interest of policymakers and researchers in the duration on uninsurance spells (e.g., Ayanian et al. 2000, Cutler & Gelber 2009). Intuitively, the fact that care can be delayed means that there is a conceptually important difference, for example, between being uninsured for six straight months and cycling in and out of insurance on a daily basis for one calendar year. Longer durations of uninsurance can lead to the accumulation of many untreated problems, and this accumulation may have severe long-term health consequences.

[55] It should be noted that this topic is conceptually related to the design of optimal income taxes in which the social planner must account for the fact that individuals can avoid taxes by re-timing income. A number of papers have found evidence of income re-timing in order to avoid taxes (e.g., Goolsbee 2000, Burman & Randolph 1994).

My analysis in this paper illustrates that this breakdown can be particularly dramatic when many insured treatments are not urgent. Within health care, treatments span the spectrum of urgency. While treatment for a heart attack is extremely urgent, knee replacements can be delayed for years. Still, all health care spending, urgent or not, is typically covered by the same insurance product, with the notable exceptions being the historical exclusion of dental and vision services. One could imagine many other ways to bundle (or unbundle) health care services for the purpose of insurance. Ex ante, it is not obvious what an optimal grouping of services would look like from an efficiency perspective. Should urgent and less urgent types of treatments be separately insured? If so, how should the design of these insurance products differ? Alternatively, if the risk of urgent care is large enough, can we obtain more efficient insurance for all treatments by bundling all care together for the purpose of insurance? I view these as important questions for future work.

Lastly, the analysis in this paper relates to the design and implementation of the recent health care reform. My findings suggest that restricting the frequency of choice through open-enrollment periods can be useful in preventing adverse selection based on pre-existing events. The recent health care legislation, the Affordable Care Act (2010), calls for the creation of health insurance exchanges through which it is expected that many Americans will buy their health insurance coverage. The act limits the use of waiting periods for insurance and eliminates most medical underwriting, potentially exposing health insurance markets to much more strategic timing.[56] Because most health care treatments can be delayed by a number of days, the analysis in this paper suggests that allowing individuals to re-evaluate insurance decisions at a high frequency can cause large inefficiencies in insurance, and in the extreme case, may cause the unraveling of the market for insurance products more generous than the minimum mandated coverage.[57] As policymakers begin to design these exchanges, my results suggest that occasional open-enrollment periods could be useful in overcoming some adverse selection associated with delaying care. The optimal frequency of such enrollment periods in health insurance (annual, biannual, etc.) is an important topic for future research.

References

Abaluck, J., and J. Gruber (2009): Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. NBER working paper #14759.

Akerlof, G. (1970): The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. Quarterly Journal of Economics, 84, 488-500.

Arrow, K. (1975): Two Notes on Inferring Long Run Behavior from Social Experiments. Publication #P-5546, Rand Corporation.

[56]The Affordable Care Act potentially exposes insurance markets to a lot of strategic timing because of two features: the elimination of excessive waiting periods and the elimination of most medical underwriting. The act does suggest the use of open-enrollment periods during implementation. A similar reform with these two features (plus no specified open-enrollment periods) has been implemented in Massachusetts, and insurers in Massachusetts have complained of such strategic selection (perhaps driven by timing) in which individuals sign up for insurance coverage for a month or two, run up abnormally large health care bills, and then drop coverage (see, e.g., Lazar 2010, Welch & Giesa 2010). In response, lawmakers in Massachusetts recently passed a law restricting enrollment to two annual periods in 2011 and just one annual period starting in 2012 (Massachusetts Legislation Chapter 288, August 10 2010).

[57]Adverse selection among insurance products of varying generosity may be severe. In addition, there is some question about whether the penalties for violating the act’s health insurance mandate will be large enough to deter adverse selection along the insurance/no insurance margin as well.

Ayanian, J., J. Weissman, E. Schneider, J. Ginsburg, and A. Zaslavsky (2000): Unmet Health Needs of Uninsured Adults in the United States. Journal of the American Medical Association, 284, 2061-2069.

Bajari, P., H. Hong, A. Khwaja, and C. Marsh (2010): Moral Hazard, Adverse Selection, and Health Expenditures: A Semiparametric Analysis. mimeo.

Buchmueller, T., and P. Feldstein (1997): The Effect of Price on Switching Among Health Plans. Journal of Health Economics, 16, 231-247.

Bundorf, K., J. Levin, and N. Mahoney (2010): Pricing and Welfare in Health Plan Choice. mimeo, Stanford University, available at http://www.stanford.edu/∼jdlevin/research.htm.

Burman, L., and W. Randolph (1994): Measuring Permanent Responses to Capital-Gains Tax Changes in Panel Data. American Economic Review, 84, 794-809.

Card, D., C. Dobkin, and N. Maestas (2008): The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare. American Economic Review, 98, 2242-2258.

Cardon, J., and I. Hendel (2001): Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey. Rand Journal of Economics, 32, 408-427.

Cardon, J., and M. Showalter (2001): An Examination of Flexible Spending Accounts. Journal of Health Economics, 20, 935-954.

Carlin, C., and R. Town (2009): Adverse Selection, Welfare, and the Optimal Pricing of Employer-Sponsored Health Plans. mimeo, University of Minnesota Working Paper.

Chiappori, P.-A., B. Jullien, B. Salanie, and F. Salanie. (2006): Asymmetric Information in Insurance: General Testable Implications. Rand Journal of Economics, 37, 783-798.

Chiappori, P.-A., and B. Salanie (2000): Testing for Asymmetric Information in Insurance Markets. Journal of Political Economy, 108, 56-78.

Cochrane, J. (1995): Time-Consistent Health Insurance. Journal of Political Economy, 103, 445-473.

Cohen, A. (2005): Asymmetric Information and Learning: Evidence from the Automobile Insurance Market. Review of Economics and Statistics, 87, 197-207.

Cohen, A., and L. Einav (2007): Estimating Risk Preferences from Deductible Choice. American Economic Review, 97, 745-788.

Cutler, D. and A. Gelber (2009): Changes in the Incidence and Duration of Periods without Insurance. New England Journal of Medicine, 360, 1740-1748.

Cutler, D., B. Lincoln, and R. Zeckhauser (2010): Selection Stories: Understanding movement across health plans. forthcoming in Journal of Health Economics.

Cutler, D., and S. Reber (1998): Paying for Health Insurance: The Trade-off Between Competition and Adverse Selection. Quarterly Journal of Economics, 113, 433-466.

Einav, L., A. Finkelstein, and M. Cullen (2010): Estimating Welfare in Insurance Markets Using Variation in Prices. Quarterly Journal of Economics, 125, 877-921.

Einav, L., A. Finkelstein, and J. Levin (2010): Beyond Testing: Empirical Models of Insurance Markets. Annual Review of Economics, 2, 311-336.

Einav, L., A. Finkelstein, and P. Schrimpf (2010): Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market. Econometrica, 78, 1031-1092.

Finkelstein, A., K. McGarry, and A. Sufi (2005): Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance. American Economic Review Papers and Proceedings, 95, 224-228.

Finkelstein, A., and J. Poterba (2004): Adverse Selection in Insurance Markets: Policyholder Evidence from U.K. Annuity Market. Journal of Political Economy, 112 Part 1, 183-208.

Gertner, R. (1993): Game Shows and Economic Behavior: Risk-Taking on "Card Sharks." Quarterly Journal of Economics, 108, 507-521.

Goda, G., and C. Manchester (2010): Incorporating Employee Heterogeneity into Default Rules for Retirement Plan Selection. mimeo, available at http://www.stanford.edu/∼gopi/PensionDefault.pdf.

Goolsbee, A. (2000): What Happens When You Tax the Rich? Evidence from Executive Compensation. Journal of Political Economy, 108, 352-378.

Guay, A. (2006): The Differences between Dental and Medical Care: Implications for Dental Benefit Plan

Design. Journal of the American Dental Association, 137, 801-806.

Handel, B. (2010): Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts. mimeo, University of California at Berkeley, available at http://www.econ.berkeley.edu/∼bhandel/index.shtml.

Hendel, I., and A. Lizzeri (2003): The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance. Quarterly Journal of Economics, 118, 299-327.

Jeffcoat, M. (2004): A Double-Edged Sword: Insurance A to Z. Journal of the American Dental Association, 135, 842-843.

Kowalski, A. (2010): Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care. mimeo, Yale University, available at http://www.econ.yale.edu/∼ak669/research.html.

Lazar, K. (2010): Short-term Insurance Buyers Drive up Cost in Massachusetts: Lawmakers Look to Close Loopholes. Boston Globe June 30, 2010.

Lustig, J. (2010): Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare. mimeo, Boston University, available at http://people.bu.edu/jlustig/research.html.

Madrian, B., and D. Shea (2001): The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. Quarterly Journal of Economics, 116, 1149-1187.

Manning, W., H. Bailit, B. Benjamin, J. Newhouse (1986): The Demand for Dental Care: Evidence from a Randomized Trial in Health Insurance. Rand publication series.

Manning, W., J. Newhouse, N. Duan, E. Keeler, A. Leibowitz and M. Marquis (1987): Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment. American Economic Review, 77, 251-277.

Metcalf, C. (1973): Making Inferences from Controlled Income Maintenance Experiments. American Economic Review, 63, 478-483.

Metrick, A. (1995): A Natural Experiment in "Jeopardy!". American Economic Review, 85, 240-253.

Palmer, C. (2009): Growth in dental spending expected to slow in 2009. American Dental Association News, accessed 10/31/2011 at http://www.ada.org/news/1103.aspx.

Rothschild, M., and J. Stiglitz. (1976): Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. Quarterly Journal of Economics, 90, 629-649.

Strombom, B., T. Buchmueller, and P. Feldstein (2002): Switching Costs, Price Sensitivity, and Health Plan Choice. Journal of Health Economics, 21, 89-116.

Sydnor, J. (2009): (Over)insuring Modest Risks. mimeo, Case Western Reserve University, available at http://faculty.weatherhead.case.edu/sydnor/.

Tchernis, R., S.-L. Normand, J. Pakes, P. Gaccione, and J. Newhouse (2006): Selection and Plan Switching Behavior. Inquiry, 43, 10-22.

Welch, D., and K. Giesa (2010): Analysis of Individual Health Coverage in Massachusetts Before and After the July 1, 2007 Merger of the Small Group and Nongroup Health Insurance Markets. Report for Massachusetts Division of Insurance, available at http://www.mass.gov/Eoca/docs/doi/Companies/adverse_selection_report.

Figure 1: Individual Out-of-Pocket Spending as a Function of Total Spending by Plan

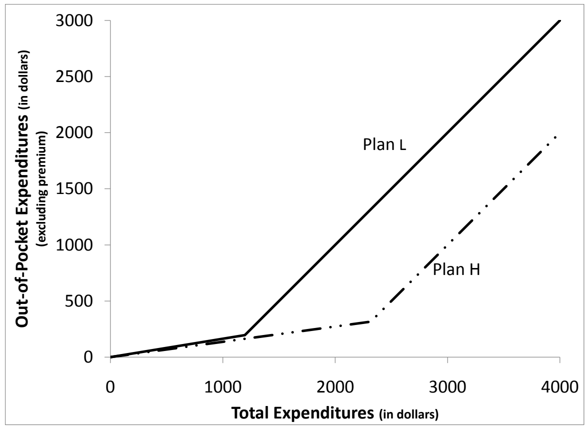

Notes: The above is a plot of annual out-of-pocket spending (excluding premiums) per individual as a function of total annual spending by plan, using the unconditional average coinsurance for the baseline sample below the annual individual maximum benefit. The details of the average coinsurance calculation are described in Table 2. Based on these out-of-pocket cost functions with the average coinsurance rate across all ages, it would take approximately $1,225 of dental spending for a single coverage employee facing $65 Plan H premium to be indifferent ex post between the two plans. The kinks in this figure are at $1,188 for Plan L, and $2,299 for Plan H (these are the levels of total spending that correspond to exhausting the $1,000 and $2,000 maximum benefits, respectively). Above these values of total dental spending, the individual pays the full cost of care.

Figure 2: Annual Individual Total Expenditures and Claim Cost for the Baseline Sample

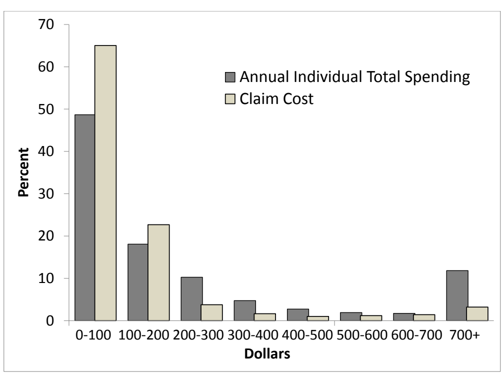

(a) Annual expenditures and claim-level costs

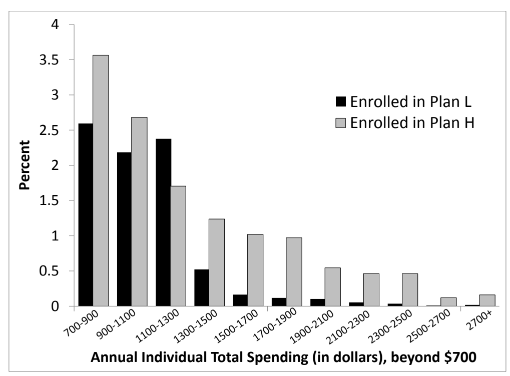

(b) Annual expenditures exceeding $700

Notes: Panel (a) above displays the distribution of claim cost (costs within one day) and annual total dental spending for the baseline sample. Observations are pooled across years to create this histogram. Thirty-eight percent of individual-years are exactly at zero dollars of annual dental spending. Panel (b) shows the distribution of annual dental spending for the right tail of the baseline sample by plan. The percent here is calculated based on the percent of individual-years out of all individual-years on the given plan (even those not described in the figure because their spending is less than $700).

Figure 3: Strategic Claim Delay

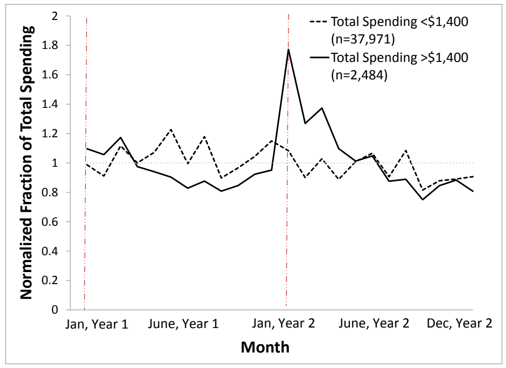

Notes: The figure displays the average normalized monthly individual dental expenditures as a fraction of the total expenditures across two adjacent years, the average value of (monthly spending $\times$ 24/total spending). The sample used to create this figure is restricted to those who were insured with the company for two adjacent years (2004-2005, 2005-2006, or 2006-2007) and were enrolled in Plan L (which has a maximum benefit of $1,000) in the first of these adjacent years. Individuals with no spending across the two years are dropped. This series is displayed separately for those with overall expenditures less than $1,400 (those who probably did not have the incentive to delay claims) and those with overall expenditures exceeding $1,400 (those who were more likely to have the incentive to delay claims). For those with high overall expenditures, one can see a spike beginning in January of the second year. For these incentivized individuals, on average 32% of the total dental spending across the two years is incurred during the first six months of year 2. Appendix A contains alternative figures with different cutoff values to identify incentivized individuals and alternative figures that control for year 2 plan choice. The qualitative patterns remain the same in these alternative figures.

Figure 4: Strategic Claim Delay When Coverage Decisions are Unlocked and Locked

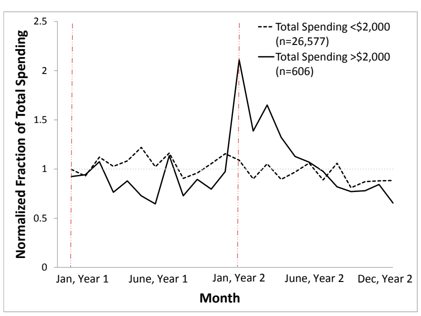

(a) Unlocked coverage

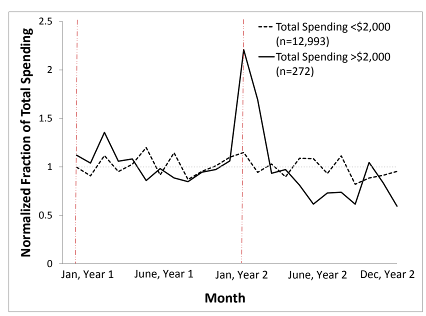

(b) Locked coverage

Notes: These figures display the average normalized monthly individual dental expenditures as a fraction of the total expenditures across two adjacent years, the average value of (monthly spending $\times$ 24/total spending). The sample used to create Panel (a) is restricted to those who were insured with the company for two adjacent years, 2004-2005 or 2006-2007, and were enrolled in Plan L during the first of these adjacent years. These individuals had the opportunity to switch to Plan H for the second of the adjacent years. The sample used to make Panel (b) is restricted to those who were insured with the company on Plan L during the locked period 2005-2006 (in particular, this panel excludes the few employees that switched coverage during the locked period due to allowable reasons like marriage or birth of a child). Individuals with no spending across the two years are dropped. In each panel, the series is displayed separately for those with overall expenditures less than $2,000 (those with less incentive to delay claims) and those with overall expenditures exceeding $2,000 (those who were more likely to have the incentive to delay claims). A higher cutoff is chosen here than in Figure 3 because the opportunity to buy more insurance should mostly affect those with very high expenditures. For those with high overall expenditures, one can see a spike in dental spending beginning in January of year 2 in both of the panels above. In the unlocked periods as compared to the locked period, on average 7% more of the total spending across the two years is done in the first six months of year 2. This pattern reflects that the annual maximum benefit (applicable in both the locked and unlocked periods) and the opportunity to switch insurance coverage (applicable in only the unlocked periods) are both important incentives for delaying treatment.

Table 1: Description of Employees and Sample Restrictions

	All Employees	Employees in Baseline Sample	Employees in Restricted Sample	US employees with dental insurance
Employee-years	181,552	116,426	34,838	
(unique employees)	(64,466)	(43,412)	(17,419)	
Male	76%	73%	72%	53%
Age (median)	44	44	36	44
Rural	39%	39%	38%	
Job Tenure (median)	7	7	11	
Union	35%	17%	>1%	7%
Wage				
Average	$47,512	$50,075	$57,126	$48,943
Median	$42,470	$42,062	$48,552	$39,000
Dental Coverage Tier				
emp single	30%	32%	22%	
emp + spouse	12%	16%	24%	
emp + children	5%	7%	9%	
emp+ family	53%	44%	44%	

Notes: In the left panel above, all the statistics are for employees, not the associated dependents. The employee-year level of observation is used when calculating the mean and median statistics. The "All Employees" column describes all employee-years for all employees that were with the company at any point between 2004 and 2007. The "Employees in Baseline Sample" column describes all employees who ever had the relevant benefit menu and were employed for the entire relevant calendar year. These employees along with their associated dependents make up the "Baseline Sample" used in Section 4 to identify evidence of strategic claim timing. The "Employees in Restricted Sample" column describes employees that are in the baseline sample and were employed from 2005-2006. These employees along with their associated dependents make up the "Restricted Sample" used in the estimation of the structural model. The "% Rural" is the percent of the sample that lives in a municipality characterized as rural by the 2000 US Census. For comparison, the "Employed US Population with Dental Insurance" column lists some descriptive statistics for the sample of people in the 2007 Medical Expenditure Panel Survey (MEPS) who were continuously employed and reported having dental insurance throughout 2007. All the values for these employees are for the year 2007. Because the MEPS does not indicate the source of dental coverage, this MEPS sample includes both employees that obtained coverage through their own employer and employees that obtained coverage from other sources (for example, through a spouse's employer). Overall, approximately 40% of the people in the MEPS 2007 report having dental insurance.

Table 2: Description of Dental Insurance Benefits by Plan

	Plan Coverage		Categories of Care[a]	
	Plan L	Plan H	% of Total Spending[b]	% of Total Claims[b]
Coinsurance Below Annual Maximum Benefit[c]				
Preventive Care[d]	100%	100%	32%	57%
Basic Care[e]	85%	85%	63%	42%
Major Care[e]	50%	50%	5%	1%
Oral Surgery[e]	50%	100%	<1%	<1%
Average Coinsurance[f]	84.2%	87.0%		
Annual Maximum Benefit	\$1,000/person	\$2,000/person		

[a] Claim categories are inferred by combining the procedure codes and claim reimbursement information. The average out-of-pocket spending to total spending ratio is calculated for each procedure code, and these codes are then classified into the care categories above. This process left less than 5% of claims with unclassifiable codes, and these claims are omitted from the statistics on the percentage of claims and spending by category. In addition to the plan differences noted above, Plan H provides orthodontia coverage for children under 18 years of age up to a separate lifetime maximum benefit of \$1,500. In the empirical analysis, I use the annual maximum benefit feature to identify strategic claim timing. Since orthodontia is not subject to this annual maximum benefit, I exclude orthodontia claims from the analysis. Only 2% of households contain an individual with an identifiable orthodontia claim in a given year. Because I cannot perfectly classify claims into those that are orthodontia and non-orthodontia, those individuals with any orthodontia claims are dropped from the analysis (or the entire associated household is dropped when doing household-level analysis).

[b] The "% of total spending" is the percent of total dental spending for each care category for individuals in the baseline sample, and the "% of total claims" is the percent of total claimed procedures for each care category for claims submitted by those in the baseline sample. The usage of the word "claims" in the heading of this table differs from the usage throughout the paper. Throughout the paper, I use "claims" to describe the total claimed procedures in a day. In contrast, here "claims" is at a more disaggregated level, as an individual may have claimed procedures that span multiple care categories above within one day.

[c] The displayed percentages are the percent of expenditures paid by the insurer for care in each of the above categories below the annual maximum benefit. Beyond the annual maximum benefit, all dental spending is the responsibility of the patient.

[d] The company places some limits on the annual amount of covered preventive cleanings and diagnostic X-rays. For example, covered patients may have up to two preventive cleanings and two partial mouth X-rays reimbursed within one calendar year.

[e] These types of care are subject to an annual individual deductible—\$50 for Plan L and \$25 for Plan H. This deductible is smaller than the vast majority of expenses in any of these categories, so one can think of it as simply a factor which increases the coinsurance rate for these services. This deductible is taken into account when calculating the average coinsurance rate.

[f] The "Average Coinsurance" rate displayed in the table is the average percent of expenditures paid by the company below the annual individual maximum benefit, where the average is taken over the different types of care accounting for the relative shares of overall dental spending in the data and accounting for the small deductible that is applicable to some care. Specifically, the formula for the average coinsurance rate can be written as follows: $\gamma_j = b_j/(\frac{b_j}{\Sigma_c w_c \gamma_{j,c}} + d_j)$. Here, j represents the plan and c represents the type of care, w_c represents the fraction of total spending in a care category c, d_j is the individual deductible for Plan j, $\gamma_{j,c}$ is the coinsurance rate for Plan j and category c, and b_j is the annual maximum benefit for Plan j. This formula is more complicated than a simple weighted average in order to adjust for the small deductible. The shares of spending in the baseline sample are used in this formula to calculate the "Average Coinsurance" rate above. In the analysis, the average coinsurance rate is conditioned on age group: less than 19 years of age, 19-30 years of age, 31-40 years of age, 41-50 years of age, and over 50 years of age.

Table 3: Descriptive Statistics

	Baseline Sample		Restricted Sample	
	Plan L	Plan H	Plan L	Plan H
% Households-years	24%	76%	24%	76%
Average % ind with zero claims	38%	38%	31%	33%
Individual dental expenditures				
Mean	$188	$257	$225	$305
Median	$99	$112	$134	$144
Std dev	$293	$420	$325	$466
% Individuals reached maximum benefit	2.1%	0.6%	3.0%	0.7%
% Individuals within $200 of maximum benefit	6.5%	0.9%	5.4%	1.2%
# Unique individuals	118,112		46,271	
#Unique households	43,412		17,461	

Notes: The "Baseline Sample" column describes all employees and dependents who ever had the relevant benefit menu and were associated with employees who remained with the company for the entire relevant calendar year. This sample is used in Section 4 to identify evidence of claim timing. The "Restricted Sample" column describes employees and dependents that are in the baseline sample and are associated with employees who were employed from 2005-2006 and selected dental coverage from the relevant menu. This sample is used in the estimation of the structural model. The "% Household-years" is the percent of household-year observations on each plan. The "Average % ind with zero claims" is the percent of individual-years with zero claims on the relevant plan. The individual dental expenditure statistics are calculated across all individual-year observations. The "% Individuals reached maximum benefit" is the percent of individual-years that exceed the level of total spending that would exhaust the maximum benefit of the plan in which they were enrolled. The "% Individuals within $200 of maximum benefit" is the percent of individual-years that would exhaust the annual individual maximum benefit of the relevant plan with $200 dollars more in total dental spending. The maximum benefit for Plan L is $1,000 and for Plan H is $2,000.

Table 4: Testing for Asymmetric Information

	(1) $\text{Claims}_{h,t}$	(2) $\text{Claims}_{h,t}$	(3) $\text{Claims}_{h,t}$	(4) $\text{Claims}_{h,t}$
$\text{Choice}_{h,t}$, defn 1	267.0***	162.1***		
	(5.16)	(9.56)		
$\text{Choice}_{h,t}$, defn 2			205.7***	128.1***
			(5.04)	(10.74)
Fixed Effects				
Premium Menu	x	x	x	x
Coverage Tier			x	x
Household		x		x
Sample Restriction			enroll same dependents	enroll same dependents
Dep Var				
Mean (Median)	598(307)	598(307)	600(308)	600(308)
Std Dev	757	757	759	759
N	116,426	116,426	100,241	100,241

Notes: OLS regression results are displayed above. The level of observation is the household-year. The sample in columns (3) and (4) is restricted to household-year observations for which the associated employee chose the same dependent coverage tier for medical and dental insurance in all years he is in the data. The dependent variable, "$\text{Claims}_{h,t}$," is the amount the company would reimburse for household h's dental expenses had the household been enrolled in Plan H during year t (regardless of the actual enrollment of the household). This amount is calculated by applying the Plan H cost-sharing rules to the different types of spending the household did. The "$\text{Choice}_{h,t}$, defn 1" variable indicates whether the employee associated with household h chose Plan H in year t and covered all his "potential dependents," where his "potential dependents" are defined as all the dependents covered by the employee for either medical or dental insurance. Included in specifications (3) and (4), the variable "$\text{Choice}_{h,t}$, defn 2" indicates whether household h enrolled in Plan H in year t. There is an additional assumption needed to calculate $\text{Claims}_{h,t}$ under the first definition of $\text{Choice}_{h,t}$: for households who leave some potential dependents uninsured, it is assumed that uninsured dependents have no claims because spending data is not observed for these uninsured dependents. All specifications include fixed effects for dental insurance premium menus, which vary slightly across employee benefit groups. Robust standard errors are clustered at the household level.
* pvalue< 0.10, ** pvalue< 0.05, *** pvalue< 0.01

Table 5: Parameter Estimates from Model

		Baseline Estimates
μ_λ		
	Constant	0.559
		(0.0119)
	Age/100	-0.4718
		(0.0265)
	$(\text{Age}/100)^2$	0.0216
		(0.0014)
	Male	-0.144
		(0.0128)
σ_λ		0.900
		(0.0124)
p_α		0.482
		(0.0485)

Notes: The parameter estimates from the empirical model are reported above, where the coefficient of absolute risk aversion is set equal to 2.3×10^{-3}. The "Restricted Sample" described in Table 3 is used to obtain these estimates. The cost intensity distribution is the empirical distribution of claim costs for the restricted sample. Estimation of β_λ, σ_λ, and p_α is done using the method of maximum simulated likelihood, taking the empirical cost intensity distribution as given. Estimation details are in Appendix B. Bootstrapped standard errors are reported above using 100 bootstrap iterations. The estimated risk distribution implies that the median number of average annual dental events is 1, and the average number of events received annually is 2. The implied 90th percentile of average annual dental events is 5. The delay probability above indicates that 48% of individuals delay treatment when incentivized to do so.

Table 6: Insurance Enrollment and Insurer Costs by Contractible Information

	Contractible Information					
	No Information		Risk Type		Pre-Existing Events	Both Risk Type and Pre-Existing Events
	Ind Case	Proc Case	Ind Case	Proc Case		
Insurance Enrollment						
No Subsidy	0.01%	0.1%	2.0%	0.6%	9.1%	100%
25% Subsidy	0.02%	6.1%	2.9%	9.7%	24.2%	100%
50% Subsidy	3.1%	21.6%	4.6%	99%	55.9%	100%
75% Subsidy	34.5%	63.4%	47.8%	100%	87.4%	100%
Average Insured Cost						
No Subsidy	37,712	6,719	5,649	4,455	880	285
25% Subsidy	30,661	1,808	4,504	1,366	632	285
50% Subsidy	4,257	886	3,258	287	430	285
75% Subsidy	771	422	588	285	319	285

Notes: In the two panels above, each cell represents an equilibrium of a different counterfactual scenario in a market for comprehensive dental insurance. "Insurance Enrollment" is the percent of individual-years insured out of the years when insurance is available. The "Mean Insured Cost" is the average cost across the individual-years insured. The two "No Information" columns represent scenarios when the insurer can price no information about the individuals. The "Risk Type" columns represent the scenarios when the insurer can price individuals' risk types, λ in the model. The "Pre-Existing Events" column represents the case when the insurer can contract on pre-existing events (delayed treatments). In this case, there are no delayed treatments in equilibrium. The "Both Risk Type & Pre-Existing Events" column represents the case in which the insurer and individual have symmetric information, and this information is contractible. The values above are for the calculated equilibrium (details in Appendix B), which is calculated separately for the individual-specific delay cost case ("Ind Case") and the procedure-specific delay cost case ("Proc Case") in the scenarios where individuals have the incentive to delay claims. When delay costs are individual-specific, it is assumed that individuals have the same delay cost for all ten years. When delay costs are procedure-specific, it is assumed that individuals independently draw a new delay cost in each year after making their insurance coverage decision for the year.

Table 7: Insurance Enrollment and Insurer Costs by Choice Frequency Restriction

	Decision Frequency						
	Individual-Specific Delay Cost			Procedure-Specific Delay Cost			
	Annual	Every 2 years	Every 5 years	Annual	Every 2 years	Every 5 years	Lifetime
Insurance Enrollment							
No Subsidy	2.0%	0.02%	0.1%	0.6%	0.1%	0.3%	100%
25% Subsidy	2.9%	0.6%	75.5%	9.7%	11.2%	82.2%	100%
50% Subsidy	4.6%	61.0%	75.7%	99%	74.5%	99%	100%
75% Subsidy	47.8%	61.0%	75.7%	100%	94.9%	100%	100%
Average Insured Cost							
No Subsidy	5,649	18,856	7,542	4,455	5,281	2,250	285
25% Subsidy	4,504	6,197	378	1,366	1,077	247	285
50% Subsidy	3,258	468	377	287	285	277	285
75% Subsidy	588	468	377	285	271	285	285

Notes: In the two panels above, each cell represents the equilibrium of a different counterfactual scenario in a market for comprehensive dental insurance. In each simulation, it is assumed that insurers can price risk type (λ), but cannot contract on pre-existing events (delayed treatments). "Insurance Enrollment" is the percent of individual-years insured out of the years when insurance is available. The "Mean Insured Cost" is the average cost across the individual-years insured. An equilibrium is calculated (details in Appendix B) for each choice frequency restriction: annual insurance selection, every 2 years, every 5 years, and lifetime. This is calculated separately for the individual-specific delay cost case and the procedure-specific delay cost case in the scenarios where individuals have the incentive to delay claims (in all but the lifetime enrollment case). When delay costs are individual-specific, it is assumed that individuals have the same delay cost for all ten years. When delay costs are procedure-specific, it is assumed that individuals independently draw a new delay cost in each year after making their insurance coverage decision for the year (if applicable).