Journal of Health Economics 61 (2018) 93–110

Contents lists available at ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/econbase

Deriving risk adjustment payment weights to maximize efficiency of health insurance markets

Timothy J. Layton[a,*], Thomas G. McGuire[b], Richard C. van Kleef[c]

[a] Department of Health Care Policy, Harvard Medical School and NBER, United States
[b] Department of Health Care Policy, Harvard Medical School and NBER, United States
[c] Erasmus School of Health Policy and Management, Erasmus University Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 6 July 2017
Received in revised form 3 July 2018
Accepted 5 July 2018
Available online 23 July 2018

Keywords:
Health insurance
Risk adjustment
Adverse selection

ABSTRACT

Risk-adjustment is critical to the functioning of regulated health insurance markets. To date, estimation and evaluation of a risk-adjustment model has been based on statistical rather than economic objective functions. We develop a framework where the objective of risk-adjustment is to minimize the efficiency loss from service-level distortions due to adverse selection, and we use the framework to develop a welfare-grounded method for estimating risk-adjustment weights. We show that when the number of risk adjustor variables exceeds the number of decisions plans make about service allocations, incentives for service-level distortion can always be eliminated via a constrained least-squares regression. When the number of plan service-level allocation decisions exceeds the number of risk-adjusters, the optimal weights can be found by an OLS regression on a straightforward transformation of the data. We illustrate this method with the data used to estimate risk-adjustment payment weights in the Netherlands (N = 16.5 million).

1. Introduction

Health insurance markets are vulnerable to market failures related to adverse selection (Einav et al., 2010a; Glazer and McGuire, 2000; Geruso and Layton, 2017). Risk adjustment of payments to health plans (aka “risk equalization”) is a widely used policy intended to counter adverse selection problems and is a fundamental component of the regulated private health insurance markets that serve as the basis of national health policy in Germany, Israel, the Netherlands, Switzerland, and other countries, as well as of key sectors in the U.S., including the Medicare Advantage program for Medicare beneficiaries and the state-level Marketplaces created by the Affordable Care Act (2010). Each of these individual health insurance markets includes a payment system, which, depending on the country, adjusts plan payments to age, gender, geographic area, past or current medical diagnoses, past spending, and other characteristics of enrollees.

To date, the payment weights attached to the different individual characteristics included in a risk adjustment model used in a given health plan payment system have been generated using regression techniques, typically via an individual-level ordinary least squares (OLS) regression of total annual health care spending on the variables included in the model (risk adjustors). The payment to the insurer for a given enrollee is then effectively set equal to the predicted value the regression model generates for that enrollee. Such a method chooses payment weights that maximize the statistical “fit” (i.e., the R-squared) of plan revenues to costs at the individual level. However, as has been pointed out in previous work, it is unclear whether a statistical measure such as the R-squared is the “correct” objective function to maximize given the goals of either the regulator or the social planner (Glazer and McGuire, 2002).

Indeed, empirical studies evaluating different risk adjustment models imply that maximizing the R-squared is not the regulator’s objective. Such studies tend to emphasize group-level fit of plan revenues to costs rather than individual-level fit. For example, Kautter et al. (2014) first estimated the federal model proposed for the U.S. Marketplaces using OLS, and then evaluated it by creating subgroups of individuals with particular characteristics and simulating average fit for each of these groups. McGuire et al. (2014) performed a similar evaluation of the Marketplace model. With data from the Netherlands, Van Kleef et al. (2016) first estimated a risk adjustment model, and then merged survey information with health claims to check fit for various groups of people, including those with low physical self-rated health status and those reporting chronic conditions. As far as we know, however, no explicit

* Corresponding author.
E-mail address: layton@hcp.med.harvard.edu (T.J. Layton).

https://doi.org/10.1016/j.jhealeco.2018.07.001
0167-6296/

underlying framework describes insurer behavior and market efficiency underlying the evaluation methods and measures used in these papers and by researchers and policymakers generally. In other words, there has been no explicit objective function for risk adjustment design.

In this paper we attempt to develop a framework to describe how insurer behavior and market efficiency relate to the risk adjustment payments and the payment weights that underlie them. We then use this framework to derive an objective function that can be used to estimate risk adjustment payment weights that produce efficient market outcomes according to our framework. We start with Glazer and McGuire (2002) which uses a model of the behavior of a profit-maximizing insurer to (1) study incentives faced by insurers to inefficiently ration certain services and (2) develop a method for estimating risk adjustment weights that neutralize these incentives when the number of services for which plans make separate decisions in terms of allocation is smaller than the number of variables in the risk adjustment model. Our key innovations are to (1) move beyond incentives and solve for the equilibrium service-level allocations insurers will offer in a symmetric competitive equilibrium under a given plan payment system and (2) extend the model to relate these (distorted) allocations (as well as the payment system that generated them) to consumer utility and social welfare. These innovations allow us to make a number of novel and important advances. First, we are able to transparently show the set of (implausible) conditions under which the R-squared is the correct objective function to be maximized by the regulator. Second, we are able to relax some (but not all) of the implausible assumptions underlying the use of the R-squared as an objective function and derive a new, more welfare-grounded objective function. Even under our new objective function implausible assumptions remain in order to make the function computationally feasible under standard data constraints. However, we still believe this objective function represents a major contribution because under our framework these assumptions are now both transparent and fewer in number than when using the R-squared, representing a first step toward truly "optimal" risk adjustment. Finally, we are able to develop simple, general, and easy-to-implement methods for deriving risk adjustment payment weights that maximize our new objective function, even in the entirely plausible, but previously unexplored, case where the number of services exceeds the number of risk adjusters.

These methods can effectively replace the conventional two-step "estimate-then-evaluate" approach, where policymakers and researchers first estimate payment weights for a given risk adjustment model using a statistical objective function and then second evaluate the weights using a different set of criteria, with a relatively simple one-step "estimate-to-maximize-the-objective" approach, where the regulator's true objective function is used to estimate the payment weights. For any risk adjustment model for which the number of risk adjustor variables exceeds the number of decisions plans make about service allocations, a simple constrained regression of healthcare spending on the risk adjustors in the model produces the payment weights that maximize the objective function. In other, typically more common cases, where the risk adjustment model includes fewer risk adjustors than services, there is typically no set of payment weights that fully eliminate incentives for service-level distortion. Under these circumstances the optimal (second-best) payment weights can be found via a standard OLS regression on a transformation of the data and the risk adjustors. Thus, while our methods are not perfect and still rely on strong assumptions, they improve on both the status quo and the more sophisticated methods developed in the academic literature (i.e., Glazer and McGuire, 2002) while maintaining the simplicity and minimal computational burden of those methods.

In addition to providing a new approach for deriving risk adjustment payment weights, our analysis turns up a fundamental issue in the economics of health plan payment. In order to construct measures of welfare loss, we need, unsurprisingly, a characterization of the efficient allocation of health care services with which to compare the equilibrium allocation. In the theoretical parts of the paper, we distinguish between efficient and equilibrium allocations, but when it comes to the empirical application, we need additional assumptions about an efficient allocation to apply our welfare metrics. Some of these assumptions are implicit in existing methods for estimating risk adjustment weights, and there is value to making them explicit. Specifically, we model our initial empirical analyses on the presently used assumption that does not distinguish between the efficient and the observed allocations. Later in the paper we propose an alternative approach to defining efficiency that makes use of researcher knowledge of areas of pre-existing distortions of health care services in the market.

Following our modeling exercise in Section 3, we use data from the Netherlands to illustrate the use of our welfare-grounded measure of payment system performance (the value of the new objective function given a set of payment weights) and to demonstrate the implementation of our new optimal payment weight estimation methods. We note that this is an illustrative demonstration and not an attempt to make any inferences about the Dutch health insurance market which involves more complexity (e.g. finer relevant service categories) than that which is captured in our data. The data for our empirical demonstration, described in Section 4, are the actual data used to estimate risk adjustment payment weights in the Netherlands, and include multiple years of information on medical care use and individual demographic and risk characteristics, on the full 16.5 million Dutch population. We replicate the payment weights used in the 138-variable risk adjustment model in place for 2015, and compare these weights, and the welfare implications of the weights, to the weights produced by our efficiency loss-minimizing approach. For estimation, we take the set of risk adjustor variables as given, using the actual risk adjustors employed in the Dutch model.[1] In Section 4 we also describe how we operationalize assumptions about the level at which plans make allocation decisions, how (expected) individual spending relates to total spending on a service, and how we interpret the data in terms of efficiency of the current system.

Empirical methods to estimate risk adjustment payment weights and results are described in Section 5 (and an associated appendix). We describe model fit, equilibrium service-level allocations, and overall welfare loss according to our framework associated with the weights generated by the current methods and the weights generated by the welfare-maximizing methods. Section 6 contains what we believe to be a promising extension suggested by our model of insurer behavior and market efficiency. As noted above, an estimation approach based on efficiency calls for an explicit statement of what is meant by efficiency and how this is manifest in the data. In Sections 4 and 5 we assume that the levels of spending observed in the data are efficient, which we show to be a key implicit assumption underlying the use of the R-squared as the objective function in the existing risk adjustment literature. In Section 6 we modify our procedure for deriving pay-

[1] A risk adjustment model involves choice of risk-adjustor variables as well as the weights to be assigned to these variables. Economic criteria, primarily "gameability" and clinical criteria, primarily "meaningfulness," are typically considered together with incremental contributions to statistical fit when selecting the risk adjustor variables. See Kautter et al. (2014) for discussion of this in the case of Marketplace risk adjustment, and Kronick and Welch (2014) and Geruso and Layton (2015) for empirical studies of "upcoding" in the case of Medicare Advantage plans. The loss functions we propose here could substitute for the use of "fit" in the decision about variables to include.