

AHRQ Quality Indicator Software Version 4.1
Technical Overview Webinar
May 14, 2010

Transcript

Facilitators:

- **John Bott, MSSW, MBA,
Contractor, Agency for Healthcare Research and Quality**
- **Jeffrey Geppert, EdM, JD,
Battelle Memorial Institute**
- **Warren Strauss, ScM,
Battelle Memorial Institute**

John Bott [AHRQ]: Hello, and welcome to today's webinar, the second part of a two-part webinar on present on admission and how it's used in the AHRQ quality indicators. My name is John Bott. I work under contract onsite with AHRQ on the quality indicators, and today's presentation will largely be conducted by Jeff Geppert with Battelle Memorial Institute on the AHRQ quality indicators as well.

A couple of announcements right at the top of the call, there was a POA paper that was sent out to the listserv — a link to it was sent out to the listserv — approximately a week ago. That's on the AHRQ QI site that you may want to download and refer to, to flush out the content that you hear today.

Also, currently on the AHRQ QI site is a link to the slides that were used in the call two days ago, and today's slides, if for some reason you may not be able to access today's webinar through the webinar function. We will go to the next slide. *[SLIDE]*

Just to put this into context, we've done a number of webinars so far here this year — two webinars on January 12th and 14th — had repeated it to provide a very, very broad overview of Version 4.1 of the AHRQ quality indicators.

There was a follow-up webinar a couple of weeks later in January to provide additional detail on a number of aspects of 4.1 that tend to be a bit more complex and require a bit more explaining, and then two days ago and today we're providing two additional webinars related to present on admission. That's an area that we tend to receive more questions about, and so this is an attempt to be as transparent as possible in regards to the quality indicators and aspects of the methodology.

Another thing that we wanted to note sooner than later is that later in the call we'll be taking questions when we get through the presentation portion. You may ask questions through the

phone line, and at that time the operator will provide instructions, but you also may ask questions typing it in online that we can read it off. That's more convenient for people. That's under the Q&A tab that you should be seeing at the top of your screen.

There is one other note on webinars. This is what we have currently planned for 2010 for webinars, and we will continue to explore the potential for additional webinars to occur yet this year. *[SLIDE]*

The next slide to get to today's agenda is that first we'll provide a brief overview of present on admission — a more expansive overview as presented two days ago — and then more detail on the POA model which will be the bulk of the discussion. The software tools provide an example and then we'll leave, as you can see by the time allotted, a large amount of time for questions.

One other note just before I turn it over to Jeff is that these webinar presentations Wednesday and Friday here will eventually be posted online with the audio, and later a transcript will also be posted for your convenience and to refer back to it, if you would like. At this time, I'd like to turn it over to Jeff Geppert.

Jeff Geppert [Battelle]: Thank you, John. *[SLIDE]*

Just by way of reorienting ourselves from our discussion on Wednesday where we talked about the basic rationale for developing this modeling approach, the need and how the QI program is responding to both changes in the availability of data — POA data in particular — and some of the research that has been done about the impact about the absence of POA data on inferences that we might want to draw from our risk-adjusted rates for our hospital quality indicators.

So remember that we were at a point of having two types of models that we were supporting. One was a model of our outcome variable — our outcome of interest — and our covariate factors related to comorbidities.

When POA data were available and we knew that the POA data improved both the outcome of interest, our ability to flag true adverse events and our comorbidities, our ability to accurately flag the things that were truly present on admission versus complications that occurred during the hospital stay, and then another model that also estimated our outcome of interest and our comorbidity covariates where POA was not available.

Our basic approach was going to be to use data where we had POA to estimate probabilities, or how likely something was to be POA in our response, in our comorbidity factors, and result in a hospital risk-adjusted rate that would be most likely given this actual and predicted POA data. *[SLIDE]*

We're going to start with our basic notation for our modeling. Before I go into the specifics, I just want to mention briefly sort of how the model was developed. The information that we're presenting today is based on a technical report that is currently under production. The main authors of that report are some folks here at Battelle — Michele Morara, who is the Lead Statistician in the development effort; Warren Strauss, who leads our analytic shop in our Columbus office; Dale Rhoda, who is also on the line with us today and also is an analyst in that group, and then Louise Ryan who was a Biostatistics Professor at Harvard and has now gone home to Australia where she still is active in our field. Those are kind of the principals involved in the development of this model. That's based on some empirical work that Dale has been doing of late related to the general problems of missing data.

The basic structure of the model is that we have our outcome variable Y , which we observe for patient- j and hospital- i . The important thing is that we observe Y on everybody. We have Y on all of the patients that we have in the population at risk for our measure. Y is one when the adverse outcome is present and zero otherwise.

Then as we talked about on Wednesday, we have a data element P, which tells us whether this case ought to be excluded because either the outcome of interest is present on admission, or there is some sort of excluding condition that was present on admission. We want to exclude this case from our denominator. The important thing here is that we don't observe P on everyone. We only observe P where we have POA data.

Then we have our vector of covariates, or Z, which as we talked about on Wednesday could be demographics; it could be severity measures, and it could be comorbidities. Again, just like Y, Z we observe on everybody. We don't make any assumption about the presence of POA.

Then we have a corresponding vector of covariates which we call "X," which again we do not observe on everyone. We only observe X for those discharges where POA data are available.

[SLIDE]

Our ultimate goal, the one that we care the most about is that for each discharge record we're trying to estimate a probability that they are going to have the outcome of interest and that's going to be conditional on that outcome not being present on admission — so where P is equal to zero — and it's going to be based on all of the covariate factors that we have in our model X. We want to estimate this probability based on P and X, where we observe P and X only on a subset of patients — subset of discharges.

Then once we have this probability, this expected value for Y, for each discharge we basically add that up over all of the discharges in the hospital and that gives us our expected rate. We're talking about in general here how we're calculating that expected rate. *[SLIDE]*

We have an assumption or form of that probability. We assume a logit form, a logistical model where our predicted probability is a function of these X covariates, and we also include a hospital sect because we want to take into account the fact that patients might not be

randomly allocated to hospitals. We want to take into account any clustering of patients with certain characteristics and certain kinds of hospitals, and so we want to estimate this model taking into account this hospital effect.

The other important model that is imbedded in this estimation approach is that we have this variable P, which indicates whether a particular discharge is present on admission. P is also going to be a function of these X covariates that we observe on each discharge, and we also assume that that takes a logit form and also has a hospital effect.

You'll just notice first that we're using basically the same model to estimate Y, as we used to estimate P. We used the same set of covariates. One could imagine that one might want to use a different set of covariates. One set of covariates might be particularly useful for predicting whether the outcome of interest occurred; another set of covariates might be of particular interest to know whether that outcome was present on admission, but for purposes of simplicity, we used basically the same model.

At the end of the presentation, I'll share with you some of the model performance information that we've calculated to give you a sense of how well that works. One could imagine coming up with a different model for P, and that's something that we'll look at as we further refine the modeling approach. *[SLIDE]*

I'm going to talk a little bit more about this later, but this hospital effect that I just mentioned, I mean, one could imagine a random effect kind of approach — which we don't use because of the characteristics of our data. We have a lot of outcome measures that were certainly not normally distributed across hospitals, and there are a lot of hospitals that have no observed events so we used a general estimating equation approach to estimate that hospital effect. We will talk more about that later. *[SLIDE]*

[00:15:00]

How are we going to estimate this model? We're going to be maximizing our likelihoods, and we're going to have two likelihood equations that we're going to be maximizing depending on the presence of the POA data. Where the POA data are available we observe X and we observe P. That's the likelihood function on your screen that we're going to maximize, and it's going to be based on sort of the joint probability of X and Z — Z that we observe on everyone and X as we're going to be predicting on some discharges. For the first one, we observe X and Z on everyone and so we know what that joint probability is.

In the second on the bottom equation is the likelihood that we're going to be maximizing where X and P are not observed. Because we don't know what X is or what P is, we have to integrate, and so the maximum likelihood is estimated by integrating. Now, we still observe Z and we still observe Y on everyone — but we only observe X and P on some. The challenge here for the purposes of our estimation is doing this integration over all of our missing data. *[SLIDE]*

Basically, the remainder of kind of what we're going to be talking about is how we do that integration. The challenge that we're faced with is that for many of our patient safety models, we have a lot of Xs. Remember, for some of the patient safety indicators the denominator might be all medical and surgical patients, or all surgical patients that have had operating room procedures. So there are very large denominators and very heterogeneous denominators that can fall into a lot of DRG categories and comorbidity categories.

The number of Xs that we have in our model and our number of covariates can be quite large. In some cases, like for pressure ulcers there are over 100 different variables, and so the integration problem becomes sort of technically infeasible. If you had 100 Xs, the number of sums that you would have to compute would be 2 to the 100th, which is a really big number with like 32 digits in it — 30 digits in it — and so we can't compute this directly.

Basically, what we're going to do is take advantage of something that we can compute directly, and so the quantity that we want on the screen here is this W conditional on data quantity. That's sort of the maximum likelihood quantity.

How likely is this data given different values of the parameters? In this case, the parameters that we're estimating are the Bayes on our model and these hospital effects. That's what we're interested in, but that's what we can't really compute directly. We're going to take advantage of the fact that the thing that we want — the quantity that we're trying to maximize — is proportional to something that we can estimate from the data and that's data conditional on W .

The basic approach is going to be calculating the thing that we can estimate, given that it's proportional to something that we want to estimate and that's our basic approach. The way that we calculate the thing that we can estimate is through this MCMC sampling. *[SLIDE]*

This is just basically a kind of formal description of kind of what I've just said. The thing that we want is this integral at the bottom integrating over the W — the data that we observe — given different values of the parameters that we want to estimate. We're going to take advantage of this proportionality, which we can compute by sampling. So if we sample, then we can take advantage of this proportionality to get the thing that we can't integrate over. *[SLIDE]*

So this is sort of the basic overview of the approach, and we'll go over some of the details in the slides that follow. There are a couple of things that we just want to highlight. There are a lot of potential problems with the data that would complicate the estimation of the models.

We could have covariates that are linearly dependent, and so the software that we've developed to estimate these models has a couple of preprocessing steps that ensure that the data will be suitable for the estimation. It makes sure that the data is formatted and it has to be sorted by hospital, because we're estimating these hospital effects. It examines all of the

columns of Z and the corresponding columns of X, and it eliminates any columns that are linearly dependent.

Then the software that does the estimation can allow for different values of P, and just to give you an example of what a different value of P would be like, one P could be that the probability, that the outcome of interest is present on admission. Another P could be that the probability that the discharge has an excluding condition was present on admission. Our ultimate P would be basically a combination of those. We don't really do a whole lot with that distinction in the way that the model is currently implemented, but that distinction is a feature of the model.

Then we're going to walk through the MCMC approach for how we actually estimate the model. We went through on Wednesday a little bit of discussion about the 2x2 tables that are used to come up with estimates of X based on Z, and so we won't go over that again today, but look on the slides on Wednesday's for a little bit more discussion about that process.

Then we go through the simulation approach to estimate our hospital — the discharge-level predictions. We're going to go over the equations on that next.

I did want to point out this last point. We actually do this in two different ways — both with and without the GEE. We're interested in sort of comparing those results, and there are some interesting results when you estimate a model sort of with and without accounting for within hospital clustering. It is informative in terms of how patients with particular kinds of characteristics tend to be in either high or low-performing hospitals. It's interesting in its own right to compare those two models.

We used the GEE where we're able to estimate the GEE, but there are some models where the GEE doesn't converge, and so we're basically still using *[unintelligible]* the simple logistic.
[SLIDE]

So in terms of the model fitting, I want to point out a couple of characteristics of this model fitting. What we're beginning with here is our sort of likelihood.

We want to estimate given the data that we observe — or given different values of the parameters — how likely is the data that we observe? That's our maximum likelihood estimator. We break this down into two pieces, into what we call a data model and a process model. The data model is concerned with the data that we observe versus the true data. That's these X s and X primes. X is the data that we observe; X prime is the true data. Those are the covariates. P is our POA indicator — the P s that we observe and the true P s where we don't observe it. That's the data model component.

Then we have this process model component which is the W — the data that we observed and our parameters. We want to know how likely those data are given different values of those parameters. That's the process model.

The bottom shows the data model. In order to estimate those models, these are sort of measurement error models — how much measurement error is in the model relative to the true data and the data that we are predicting. These models are based on distributions. There's something called the Dirac delta distribution, or X s that are continuous and Kronecker delta for components of X that are discrete. Those are the distributions that we use in our data model portion of the likelihood. *[SLIDE]*

Then there is the process model component, which is our W conditional on our parameters. Similarly, we break that down into pieces. We have our likelihood associated with our Y ; our outcomes; our P s, which are POA flagged; our X s which are covariates that is based on POA that we predict based on our Z s, and then we have our Z s that we observe on everybody. We observe Z on everybody.

We break this down into its components, and then we're going to estimate these different components. This is where we start to make a little bit of simplifying assumptions in order to do this. *[SLIDE]*

This shows the model that we're using that we're estimating for each one of these likelihoods, and so this goes back to those original likelihood equations that we showed in the earlier slides — with the addition so we have our PI which is our Y prediction; our IR which is our P prediction, and then we have this additional logit which is predicting our Xs based on our Zs.

[00:30:00]

But in order to be able to estimate all of these things, again we're faced with sort of a feasibility challenge, and so we start to make some simplifying assumptions. To simplify the model and make it feasible even though it runs a long time now, it would run even longer if we estimated this full model. *[SLIDE]*

There are three sort of simplifying assumptions that we make. The first simplification is the component of the process model related to R and S, and so those are the models for P or POA, and our Xs. So we fit those models once based on the data that we observe, and so basically based on the discharges that have no missing data. We just estimate those models once, and then we consider them fixed for the duration of the MCMC simulation. We estimate those values — the P and the Xs, or the Rs and the Ss — do that once and then we consider them fixed.

The second simplification that we use in order to make this feasible is based on how we sample from the data. We use something called Gibbs sampling, which as I understand it, is a special case of something more general called the Metropolis-Hasting sampling approach.

Basically, what this allows us to do is to be more efficient in how we sample from the data, and so we sort of converge on our eventual solution more quickly than we would in the absence of these sampling approaches.

The third simplifying assumption that we make, we've mentioned earlier, which was rather than doing the hospital-specific random effects we use the GEE — general estimating equation theory — in order to account for the within-hospital correlation. *[SLIDE]*

So this is how the Gibbs sampling approach is implemented. Basically, what we're doing is we're estimating our parameters — our Bayes and our hospital effects — and then we're drawing a new set of parameters for our Bayes and our hospital effects according to this normal function, which is on the right-hand side of this equation. That's our chain of estimators that we're using in our likelihood estimation. *[SLIDE]*

Then this is how we're implementing the third simplifying assumption to use the GEE that we're estimating this variance empirically, rather than using a classical model variance. *[SLIDE]*

This is sort of the summary of the modeling approach. As we use this sampling methodology, this Metropolis-Hasting sampling method for our estimation for the P and X variables that we don't observe on everyone, but then we use this Gibbs sampler for drawing our Bayes on our model of our Y and we use the GEE in that estimation. *[SLIDE]*

As I mentioned at the beginning, there are a couple of other methods that are used to make sure that the data are going to work in our estimation. We look for linear dependence on the different covariates, and this is the decomposition approach that we used to do that.

Then the other issue that we're concerned about is what's called separation. It's basically when some set of the covariates are associated with positive outcomes, and some values on

the covariates are associated with negative outcomes. That can result in some infinite estimates, and so we use this ridge regression technique to account for separability. [SLIDE]

This is sort of the rationale for that, accounting for this separability. It doesn't affect our Bayes, but it makes the solution more stable. [SLIDE]

So how is this model implemented? It's implemented in a piece of software. We estimate these models using the reference population from the state inpatient databases. For the most recent version of the software, we're using the 2007 state inpatient databases to estimate these models.

There are actually four versions of the model that are estimated, both with and without GEE. The first one is just what we think of as the traditional model which is Y based on Zs, where Y and Z are observed on everyone. Then there is the model that estimates P given the Xs, and then a model based on Y given the Xs where P is equal to zero and then the MCMC version of that.

For those of you who have looked at this software, there is a regression analysis file that has the coefficients for each one of these models so you can look at the coefficients and see how the models are estimated. The thing that I wanted to mention about that is that the parameter estimates to compute sort of the traditional logistic model — the Y given Z — are also included in those files.

For those who are interested in sort of comparing the traditional model with the modeling approach that we're describing here, you can look at those Y given Z parameter estimates and apply them to your data to see what the more traditional logistic would result in. [SLIDE]

This is just a little bit more detail about how this software works. In both the SAS and WINK version that's going to be released, underlying there is a C++ program that's called to implement this MCMC simulation.

What the SAS code and the Windows software do is they output a comma delimited file that has all the Ys, the Ps, Xs and the Zs, and then they call this C++ program. It reads in that comma delimited file and writes out a comma delimited file — and has all of the predictions for each discharge record eliminating the zero and linear-dependent columns and forming the GEE regression. Then it outputs the result, and that output file is then read in by SAS and read in by WINK in the Windows program and used in the computation of the expected rate.

[SLIDE]

There is some more of the output that is generated by the software. In those regression analysis files that come with the software that have those covariates in them, there are a couple of diagnostic information. There is the model standard errors and the empirical standard errors that are calculated, and they're in those files to give you a sense of how much variability there is in those parameter estimates. *[SLIDE]*

This is some of the additional output that the software generates. We don't use all of it in the AHRQ software, because we focus on the MCMC estimates. If you look in those prediction files there are other predictions. All of those models that I've just described are included in those prediction files, which you can look at in order to see how these different predictions differ. *[SLIDES]*

I just wanted to show you a little bit of data. This is the example that we looked at on Wednesday with postoperative sepsis where we have our Y, which is the tpps13 data element, which is something that we observe on everyone. Then we have our P, which is the qpps13 data element which is something that we only observe on a subset of people. *[SLIDE]*

These are the hospital-level estimates kind of using the traditional model, the Y given Z. Then the sort of improved version of the model — the Y given X where P is equal to zero — so this is kind of what we want as our gold standard. This is our estimates of Y where we

have P data. You can see that the R is about 0.7826, and so they're highly correlated but they're not perfectly correlated. *[SLIDE]*

Then this is sort of our gold standard model again — the Y given X where P is equal to zero — compared to our imputed or our predicted version of the model where the Ps and the Xs are predicted based on the model. You can see that they're much more highly correlated.

Again, the ultimate goal of this whole modeling approach is to reduce the bias associated with the absence of the POA data. I'm going to stop and turn it back over to John just to see. *[SLIDE]* While we're queuing up some questions, I'm going to see about maybe accessing a little bit of data to show you. John — ?

John Bott [AHRQ]: Thank you very much, Jeff, for that more in-depth overview of the present on admission that was discussed on Wednesday, and Wednesday was a bit broader. People can ask a question through the online option, which is again clicking on the Q&A and typing in your question. At this time I'll turn it over to the operator to provide instructions for how to go about asking a question over the phone. Operator — ?

[00:45:00]

[Operator]: If you would like to ask a question, please press “*1.” To withdraw your question, please press “*2.”

John Bott [AHRQ]: I'm going to forward just a couple of slides while we're waiting for questions to queue up. This is the last slide in the slide deck. I just really wanted to point out here where the Present on Admission White Paper is that we noted at the top of the call on the AHRQ website to facilitate people in being able to find that. It's about 12 or 14 pages, and it walks you through in somewhat of the format that we've used over the last couple of days — starting at a higher level and providing the methodology and going into much greater detail in

the two appendices with an example. We'll go back to where we were. At this time we have yet to have any questions.

Jeff Geppert [Battelle]: I have a couple of things to show folks, if that's all right.

John Bott [AHRQ]: Sure. You're going to display live at this time now?

Jeff Geppert [Battelle]: Yes, and so just in terms of some next steps with the model, as I mentioned at the beginning, we have kind of a common single model that we use for both the Ys and the Ps. One of the things is to continue to refine those models to come up with risk factors that are the most useful in terms of predicting Y and P. We could potentially use a different model for P than we use for Y. There might be things that are particularly useful for predicting whether something is present on admission.

In particular, there might be other triggers in the data that we don't necessarily want to include in the indicator specification, but that might suggest to us that this particular adverse event is more likely to be on present on admission. We might want to include those types of risk factors in our model. As we talked about Wednesday, this POA is an important component of improving the indicators and their sensitivity and specificity, but it only gets you so far. There are other issues related to specificity, other than POA, that we continue to look at in our validation work.

The other aspect of this, of course, is the accuracy of the POA coding. We know that we'll be getting more and more POA data as more states collect it and more hospitals collect it. The quality of that coding is still under scrutiny. As people look at this, we will be developing tools to help hospitals to assess the quality of their coding and to develop best practices. That will have another impact on the ability of the model to estimate these quantities over time.

Why don't we start again with the questions, and then I'll continue with a little bit of data after we address them.

John Bott [AHRQ]: The first question typed in is, “Can you repeat the names of the distributions assumed for the delta in Slide 13?”

Jeff Geppert [Battelle]: These are basically distributions that are used for measurement error. The two distributions are the Dirac delta. That’s d-i-r-a-c delta distribution. Those are for continuous variables, and then Kronecker delta for discrete variables. That’s k-r-o-n-e-c-k-e-r.

John Bott [AHRQ]: We have one more question typed in at this time. The person asks, “When will the modified PSI 3 model data be available?”

Jeff Geppert [Battelle]: The issue with PSI #3 — PSI #3 is pressure ulcer — is that in October of 2009 there were some new codes introduced. I may be off by a year — 2008 — but anyway there were new codes introduced for pressure ulcer staging. I’m off by a year. It’s October of 2008. Currently, the pressure ulcer indicator is defined based on site. Pressure ulcer is based on site, and so the staging coding is in addition to those site codes. Now, you code the site of the pressure ulcer and then also what stage it is in — one, two, three, four.

We’re more interested from a patient safety perspective in the higher stage pressure ulcers — 3, 4. In this version of the QIs, we modified the exclusion criteria to exclude Stage 1 or Stage 2; however, the data that we used for our reference population which is the state inpatient databases from AHRQ, don’t have that staging coding in them and so we can’t estimate our models until we get data that has the staging coding in them. There are models that are estimated, but the prevalence isn’t right because we’re estimating those models using data that don’t have staging.

The 2008, which is the data that is just becoming available, does have a quarter’s worth of those stage data. We’ll be able to estimate models that have the staging in them, but that won’t be available until the next release and so it will be a little bit of time before those

models are available. Our recommendation until those models are available has been to focus mostly on the observed rates that hospitals are able to compute on their own data.

John Bott [AHRQ]: Another question typed in, the person asks, “For the 2007 data used in the SID were there any audits done to evaluate the POA coding accuracy?” Jeff — ?

Jeff Geppert [Battelle]: Only to exclude — we talked a little bit about this on Wednesday, but the most sort of egregious instances where it looked like for a particular hospital all of the values were yes, or all of the values were no, and there didn’t seem to be any real variability into the coding where you suspected that something is sort of getting populated automatically. There is a lot of work in this area.

There are coding screens that have been discussed at the AHRQ annual meeting that were developed as part of a pilot project that AHRQ has been sponsoring looking at the POA coding. There are projects that 3M and CMS and a lot of other places are developing good, clinically sound and validated algorithms for screening POA codes. One of the things that we’re working on as we speak is assessing those screening algorithms that people are developing and looking at the HCUP data, and then assessing the impact of that on our models.

John Bott [AHRQ]: I’ll read in one more typed in question before going back to the operator. The person asks, “I’m working with older HCUP data — 1998 to 2004 — where POA indicators are generally not available to compare quality indicators across different hospitals. Are there any systematic biases that I should be aware of in terms of POA occurrence by hospital size, teaching status, rural, et cetera? Can I assume that a lack of POA information creates basically random errors across hospitals?”

Jeff Geppert [Battelle]: The way that we’ve estimated these models is that we tried to account for any variability in POA coding that can be attributed to hospitals. The intent is that when the models estimate how likely it is that a particular adverse event is POA, they

would have accounted for that kind of within-hospital variability. That's a good reason for using the models, in order to account for that. If you didn't, then you might be concerned.

As we start to generate some documentation around these models, then I think that we can get a better sense of what the data show in terms of whether there are any important hospital characteristics that seem to be associated with varying levels of POA coding.

John Bott [AHRQ]: We'll go with the next typed in question. A person asks, "What are the maximum number of diagnosis codes used in the risk adjustment model?"

Jeff Geppert [Battelle]: When we estimate the models on the reference population, we use all of the available diagnostic information. The maximum number of codes that are available is 30. When hospitals execute the software and apply those risk adjustment parameters to their own data, they're able to use as many diagnostic codes as they have available to them.

There is a little bit of a difference between the SAS and the Windows in that regard. In the Windows and the SAS there is a numeric parameter that users can set — and there is no limit to that parameter. You can set it to any number that you want — even more than 30. The Windows has a finite number of data elements that you can map to, and I think that there are 35 of those; although, some of them are intended to be reserved for e-codes.

The only other exception to that is the Limited License Grouper, so the software comes with a Limited License APR-DRG Grouper. The maximum number of diagnosis codes for that grouper is hardcoded at 35.

John Bott [AHRQ]: We have another online question. The person asks, "Have you considered creating a SAS procedure from the external module?"

Jeff Geppert [Battelle]: Basically, you're writing the external module in SAS. No, and the reason that we don't is because in order to do this efficiently, there are literally millions of calculations that are occurring kind of behind-the-scenes.

In order to do this efficiently we have to use some pretty technical methods, and basically the methods that we've come up with to do these kinds of estimations in an efficient way we'd have to rely upon that library of techniques that have been developed at Battelle in order to make this estimation possible. If we used something that was more off-the-shelf, the feasibility of it, you just couldn't do it in a reasonable amount of time.

[01:00:00]

John Bott [AHRQ]: *[Prompt operator for questions]* That's currently all of the questions that we have. Jeff, do you have a couple of other additional things to note?

Jeff Geppert [Battelle]: Yes, I wanted to just give people a little bit of sense for these models. Can you see that, John?

John Bott [AHRQ]: Yes, I can see it. The PSI #12?

Jeff Geppert [Battelle]: Yes, and I think that a lot of the questions were around how well can you really sort of predict the likelihood that a particular adverse event is present on admission. I just wanted to give some sense of what that data looks like for one particular indicator. This basically shows the hospital-level rate for present on admission — for hospitals that collect POA data.

You can see that at the hospital level it ranges from something a little bit less than three per thousand to something about 1.5 per thousand. That's the actual data — the blue line — and then the pink line is the model and how well the model predicts P based on the values of X. There is some variability around the model, but in general it tracks pretty well.

It computes some c-statistics around those models, and so in particular focusing on the column that is the P conditional on X model. That's the model that predicts how likely something is present on admission given X. All of the c-statistics for most of the measures are around .8 or a little bit higher, and so actually the models do a pretty good job of predicting Ps based on Xs.

That is sort of an important prerequisite to the ability of the model to account for bias associated with missing P. The only real exception to that is postoperative hip fracture, but that is an extremely rare outcome, and a high proportion of those cases are in fact present on admission.

In general, the models have higher c-statistics on the full models of Y and Z, than they do once you start eliminating outcomes that are present on admission. You start dropping outcomes that are POA, and then obviously your models become less predictive because you're losing a lot of your outcomes — but the MCMC approach helps to compensate for that to some extent. I just wanted to give people a little bit of a sense for empirically how the models were performing.

John Bott [AHRQ]: Yes, that was nice to see that performance on additional measures, in addition to the one that you showed on postoperative sepsis, I believe. A person asks online, “When will the MCMC module be available for UNIX systems?”

Jeff Geppert [Battelle]: It's an important question for a lot of our users. It's a technically complicated solution, and so it's something that we're working with our colleagues at AHRQ to make happen, because I know that it's important for a lot of users, but there is no fixed timetable at this point to have the UNIX version available.

John Bott [AHRQ]: There is one other question online. Could you repeat what you said about the number of diagnosis codes? I think that the person asked about that is used in risk adjustment, because I think that's what it's referring to.

Jeff Geppert [Battelle]: So there are sort of three different answers to that question. The number of diagnosis codes that are used on the reference population when we estimate the models, and the answer there is that we use all of the codes that are available on the HCUP SID data — which the maximum is 30 diagnosis codes. That's what is used in the estimation.

Then the second aspect of that question is how many diagnosis codes can hospitals use when they apply the software to their data. The answer to that is in SAS there is no limit; you can use as many codes as you want. It's a parameter that you can set. In Windows there is a fixed number to the number of diagnosis codes that can be mapped into the input data file, and I believe that the number is 35.

Then the final aspect to that question is that for those who are using the APR-DRG Limited License Grouper, how many diagnosis codes does the grouper allow, and that number is also fixed at 35.

John Bott [AHRQ]: There is one other question that came in online. The person asks, "Can you comment on the recommended number of iterations for the MCMC simulation and whether you would recommend fewer for particularly large datasets?"

Jeff Geppert [Battelle]: Yes, right now the number of iterations is fixed in the software; it's one thousand iterations. Our current recommendation is not to reduce that number, because in order to know whether or not that number can be reduced would require some data analysis and understanding how stable those estimates are. Without good guidance based on that data analysis, we'd have no particular recommendation as to how that number could be lowered. I mean, the reason you'd lower it is obviously to make the estimation run faster on large data, but right now our recommendation is that you don't do that.

John Bott [AHRQ]: That's the remaining questions typed in online. *[Prompt operator for questions]* Okay, we have gone through all of the online questions. We have no questions over the phone. At this time, I'll say thank you for calling in and participating in today's webinar. We hope that this has helped people in understanding this aspect of the methodology with the QIs. Please periodically check back for the audio portion of this webinar and the transcription that will be posted on the AHRQ QI website. At this time, again thank you for participating.

[WEBINAR CONCLUDES]