

## A FRAMEWORK FOR EVIDENCE-BASED HEALTH CARE INCENTIVES SIMULATION

Joseph P. Bigus  
Ching-Hua Chen-Ritzo

IBM T.J. Watson Research Center  
1101 Kitchawan Road  
Yorktown Heights, NY 10598, USA

Robert Sorrentino

IBM T.J. Watson Research Center  
19 Skyline Drive  
Hawthorne, NY 10532, USA

### ABSTRACT

We present a general simulation framework designed for modeling incentives in a health care delivery system. This first version of the framework focuses on representing provider incentives. Key framework components are described in detail, and we provide an overview of how data-driven analytic methods can be integrated with this framework to enable evidence-based simulation. The software implementation of a simple simulation model based on this framework is also presented.

### 1 INTRODUCTION

The health care industry is a complex system involving payers, providers, and patients. This system is currently in a state of flux as government and employers try to reduce costs while simultaneously maintaining or improving health outcomes. The proposed changes impact both payers and providers through use of alternative payment models, new provider business organization models, and use of evidence-based treatment plans. Alternatives to the current fee-for-service payment model are now being evaluated, including bundled episode-based approaches, outcome-based approaches based on improved health outcomes, and gain-sharing approaches, where physician driven cost savings are shared between the payer and the provider organization.

Due to the complexity of health care delivery, it is not easy to immediately determine the impact of using those different models on the financial and health outcomes of the healthcare system. The primary approach to exploring changes in payment fee structures have been in-field experiments or pilots focused on certain patient segments or specific provider groups. These pilots usually take much time and effort and may not be generalizable across geographical regions, patient groups and providers. In this paper, we describe a general framework for simulating the impact of alternative healthcare incentives. In addition, we discuss opportunities for using client data to enable evidence-based customization of the patient, provider, disease and intervention models used in a simulation.

Figure 1 illustrates the major components of our framework. The three primary groups of models are the *sub-systems*, the *decision-makers*, and the *evidence-based analytics*. The sub-system models describe various aspects of the clinical and socio-economic environment, and includes disease models, provider reimbursement mechanisms, and regulations or guidelines for care. The decision-makers include key players such as the patient, provider and payer. The evidence-based analytics provide the statistical, optimization, machine-learning and visualization models for aiding in model building, model calibration and result interpretation. The evidence-based analytics provide the means for customizing a simulation model for a particular population of patients and providers based on observational data available in the form of medical claims and patient health records.

One of the challenges associated with integrating several models into a single simulation has to do with integrating models that may have been developed independently using different modeling paradigms, software languages and applications. While it is outside the scope of our work to address this challenge, we note that Cefkin et al. 2010 have developed the Smarter Planet Platform for Analysis and Simulation of

Health (Splash), which is a framework that supports integration of multiple existing models, simulations, and data that represent parts of the broader health ecosystem.

The rest of this paper is organized as follows: In §2, we provide a brief review of the literature. In §3, we provide detailed descriptions of the framework components shown in Figure 1. In §4 we describe how we have implemented a simple simulation using our framework. We conclude with a discussion on directions for future work.

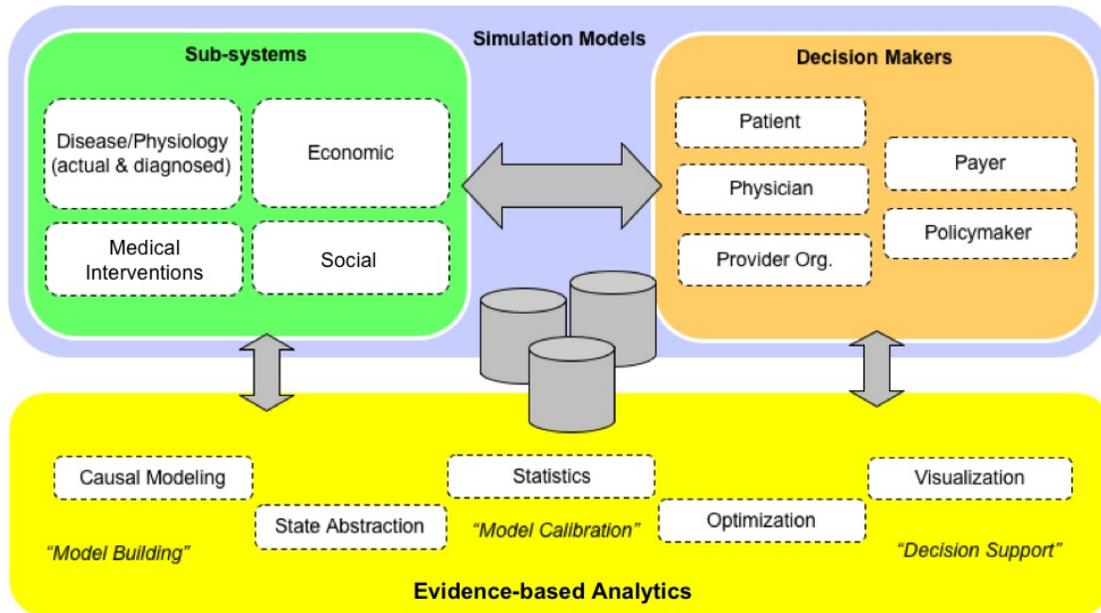


Figure 1: Conceptual Overview of Simulation Framework.

## 2 RELATED LITERATURE

The relevance of agency theory to the health care industry has been noted in the literature (Dranove and White 1989), as well as its specific applicability to physician reimbursement contracts (McGuire 2000, Fuloria and Zenios 2001, Yaesoubi and Roberts 2011). These papers, however, generally do not use results from analyses of empirical data to formulate or validate modeling assumptions. A notable exception is the recent work by Lee and Zenios, who use a data-driven approach for designing reimbursement contracts in the context of care delivery for medicare patients with end-stage kidney disease. This evidence-based approach lends credibility and realism into previously abstract models. The objective of this work is to make evidence-based approaches for studying incentives in health care available to researchers, practitioners and health policymakers by way of a convenient, flexible, and scalable simulation software platform.

Montreuil et al. 2007 developed an agent-based simulation framework for representing a health care delivery system, and later, Charfeddine and Montreuil 2010 utilized this framework to simulate a health care delivery network for patients with chronic obstructive pulmonary disease in Canada. Christiansen and Campbell 2003 have developed a comprehensive, agent-based simulation tool. Our simulation framework and platform provides a specific focus on simulating incentives in health care delivery. As a result, our framework may de-emphasize or exclude certain aspects of the health care system that we believe may contribute less to understanding the impact of these incentives.

### 3 SIMULATION FRAMEWORK

In this section we describe the key elements of our currently implemented framework. Since our initial interest has been in simulating provider responses to incentives, the current framework emphasizes the role of the physician. Future developments will be discussed in Section 5.

#### 3.1 Payers, Providers and Patients

Our framework assumes three key decision makers: the payer, the provider and the patient. In this section, we describe how they are represented in our framework.

##### 3.1.1 Payers

In our framework, a payer object (`Payer`) is defined by its member patients and network providers, and by the contracts that it holds with them. Whenever a patient receives a service from a provider, these contracts determine the financial exchanges with the payer. For example, patients pay premiums to payers and payers reimburse providers for covered services. In our framework, a `Payer` also contains an object called a `ClaimsPayoutProcess` which is used to adjudicate claims and transfer payments to providers. Claims adjudication is described in Section 3.4.

If we consider a typical problem that this framework may be used to address, the goal of a payer would be to minimize medical costs (i.e., reimbursements to providers) over the long term. This means that they would prefer to offer providers the lowest level of reimbursement that will still provide enough incentive to the providers to ensure a high enough quality of care to the payer's members (i.e., the patients) so that long-term medical costs can remain low. The specifics of how quality is measured will depend on the problem application.

##### 3.1.2 Providers

Each provider has a *panel* of patients that they care for. For each patient encounter, the provider must evaluate the patient and decide on a treatment plan. Providers typically have contractual relationships with several payers and are subject to those contract terms and conditions when treating the payer's members. A provider object (`Provider`) in our framework is characterized by various attributes, such as their professional expertise (e.g., general practitioner, endocrinologist, registered dietitian, nurse practitioner), facility affiliations and contracts with payers. We organize a provider's attributes into four categories: demographic, professional, behavioral, and economic. Examples of these attributes are summarized in Table 1.

The goal of the provider is a complex combination of altruism and personal interest. Our framework was designed in a way that emphasizes the provider's interest in optimizing their financial rewards. Therefore, in §3.4 we provide details on how financial incentives can be represented in a simulation model. Given a reimbursement contract from a payer, the provider's goal is to select medical interventions that will maximize their reimbursements. The contract between the payer and provider is used to align the provider and the payer's interests. Including certain factors in a model, such as capacity constraints and minimal standards of care, may help prevent extreme levels of care.

From our perspective, the *behavioral* category is important because it defines how the provider will select the treatments applied to a patient. When a provider selects a treatment, she has access to the payer contract that will determine her reimbursement for any services rendered. The decision model implemented here could be a static or dynamic treatment policy derived by way of a mathematical optimization or other method. A provider's behavioral models may be further parameterized by attributes in the other categories. For example, the years of experience of a provider could be predictive of the quality of their treatment decisions (Choudhry et al. 2005). As another example, a provider's gender may be predictive of their level of implementation of preventive services (Bertakis et al. 1995). A provider's specialty could also

Table 1: Summary of Provider Attributes.

Category	Description	Examples
Demographic	Provider's cultural and demographic background.	Age, Gender, Race
Professional	Provider's professional background.	Employer, Specialty, Experience, Education, Facility Affiliations
Behavioral	Decision model(s) for selecting patient treatment(s).	Various algorithms.
Economic	Financial parameters.	Wealth, Costs, Revenue, Risk Aversion, Provider Contract(s)

pre-dispose them to recommend certain types of treatment (e.g., a preference for surgery vs. non-surgical treatment). Our representation of provider interventions is presented in Section 3.3.

### 3.1.3 Patients

Patients are the agents that drive demand for healthcare services. The primary aspects of patients are their demographic and biometric attributes, their general health status, known disease states and current treatment plans. The major behavioral aspects of patients include their lifestyle habits such as smoking, diet, exercise, compliance with treatment programs and social support networks. Examples of a patient's attributes are categorized and summarized in Table 2.

The patient's goal is to optimize his quality of life, presumably by maintaining certain level of health. He does this by choosing whether to adhere to the interventions recommended by the provider, and what kind of lifestyle to adopt. Ideally, contracts between the patient and providers or between the patients and payer could be used to align the interests of patients with these other parties. Our current framework does not capture such contracts. To some extent, however, the payer's goal of keeping long-term medical costs low captures the desire of the patient to remain healthy.

Table 2: Summary of Patient Attributes.

Category	Description	Examples
Demographic & Lifestyle	Patient's cultural, social, and economic background.	Age, Gender, Race, Occupation, Education, Smoker, Alcoholic, Social Support Network, Compliance with Treatments
Biomarkers	Measurable, biological indicators.	Blood Pressure, Cholesterol, Blood Glucose, Serum Electrolytes
Diseases	Disease processes applied to the patient.	Chronic Obstructive Pulmonary Disease, Colon Cancer, Diabetes
Treatments	Therapeutic processes applied to the patient.	Diet & Exercise Program, Antibiotics, Oral hypoglycemics, Chemotherapy
Health Status	User-defined patient state. Typically, a function of the above attributes.	Controlled Geriatric Diabetic w/o comorbidities on oral hypoglycemic medications

Our framework defines two versions of each attribute in the biomarkers, diseases and health status, categories: an *actual value* and a *believed value*. The reason for this is that while the patient's clinical progression is determined by his true clinical condition, the provider's decision regarding treatment is based upon their belief about the patient's clinical condition, which is only partially observable. For various reasons, measurements of patient biomarkers may also be of limited accuracy. Each patient may have one or more associated disease models that may respond to the treatment plans applied. The representation of patient behavior in our current framework is limited to the specification of *compliance probabilities* (which may be alternatively interpreted as *consent probabilities*) for each treatment that is prescribed for the patient's by his provider. These attributes fall under the 'Demographic & Lifestyle' category in Table 2.

### 3.2 Disease Progression

A patient's health progression is a central dynamic element in health care delivery, and is a key driver of the demand for provider services. The ability to control a patient's disease progression through alternative treatments is an essential requirement for the design of effective provider incentives. In general, any disease model (`DiseaseModel`) implemented in our framework must be designed to receive a `Patient` object and modify its disease-related attributes based on any disease changes that may occur during a given time-interval. Alternatively, the disease model should determine the next disease event and the time of its occurrence. While a Markov model is a popular approach for modeling disease progression (Aalen et al. 1997, Silverstein et al. 1999, Barhak et al. 2010), our framework does not require that a specific methodology be used, as long as its instantiation implements our disease model interface, `IDiseaseModel`. To date, we have successfully implemented the model of Barhak et al. 2010 using this interface into our proof-of-concept, which is presented in Section 4.

Since each `Patient` may have one or more disease models associated with it, and each disease model may be implemented separately, our framework provides a disease model *manager* (`DiseaseModelManager`), which is collection of one or more disease models. Additionally, we provide a disease model *service* (`DiseaseModelService`), which will define the logic for coordinating the progression of multiple disease models (i.e., to the next time step, or disease event).

### 3.3 Medical Intervention

Medical interventions refer to the treatments plans that are prescribed for a patient by a provider. Since the objective of our simulation framework is to enable the simulation of a synthetic stream of claims, we provide a structure that allows every treatment plan to be decomposed into billable procedures. Hence, medical interventions serve as the sources of new claims.

#### 3.3.1 Treatments

We represent treatment plans hierarchically. At the base of the hierarchy are *provider actions*. Examples of provider actions include billable procedures, lab tests, and referrals to other providers. In general, some of the physician actions at the base of this structure could also be non-billable, such as asking a patient about their family history, or stress levels at work or home. These actions may not be billable, but could influence the physicians beliefs about the patient, and hence treatment choice and eventual health outcome. The next level up in the interventions hierarchy are treatment *plans*, which are collections of provider actions. These actions may be sequenced so that a time-series of different actions are specified by a plan. A treatment plan is stochastic in the sense that a provider will implement any of the actions comprising the plan with some probability. A provider behavioral model defines the values of these probabilities. Multiple care plans may be combined into a treatment *package*. Figure 2 illustrates the described treatment hierarchy. From Figure 2, note that one can specify time lags between actions, and probabilities of performing an action. The final

probability of an action being performed is a product of the both the provider’s probability of implementing or recommending an action, and the patient’s probability of consenting to, or complying with, the action.

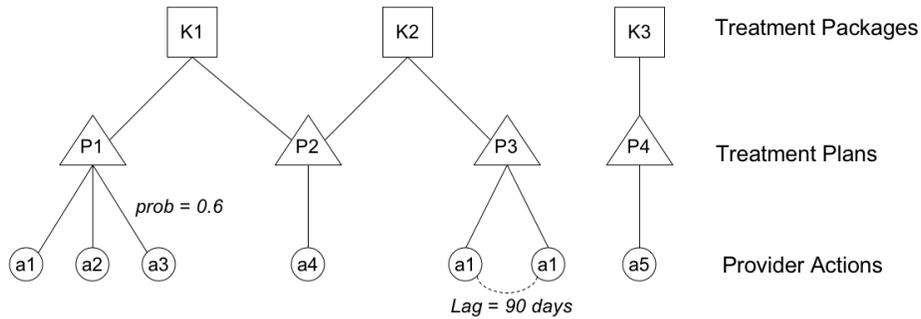


Figure 2: Medical Intervention Hierarchies.

Table 3 provides examples of provider actions, treatment plans and treatment packages corresponding to Figure 2. This example is meant to be illustrative of the framework only, and is not intended to represent recommended medical guidelines for diabetic care.

Table 3: Examples of Provider Actions, Treatment Plans and Treatment Packages.

Type	Label	Description
Action	a1	(CPT Code: 83036) Hemoglobin glycosylated (A1c) Test
	a2	(CPT Code: 99214) Office or other outpatient visit, 25 minutes
	a3	(CPT Code: G0108) Diabetes out-patient self management training services
	a4	(CPT Code: 97803) Medical Nutrition Therapy, individual, 15 minutes
	a5	(CPT Code: 67228) Panretinal photocoagulation
Plan	P1	Monitoring & Self-Management Training
	P2	Diet & Exercise Education
	P3	HbA1c Control
	P4	Laser Treatment
Package	K1	(Elective) Diabetes Care with Self-Management
	K2	(Elective) Diabetes Care, Controlled Diabetics
	K3	(Non-elective) Macular Edema Emergency

In general, treatment packages can be labeled as *elective* or *non-elective*, meaning that they are implemented in an urgent (non-elective) or non-urgent (elective) fashion. Thus, in our framework, each provider specialty possesses a ‘tool-kit’ of treatment packages that they are qualified to implement. A provider’s behavior model defines the policy for selecting a particular treatment package for a patient with certain attributes. This policy could be stochastic. In general, elective treatment packages are implemented according to the provider’s behavior model and based on patient attributes. Non-elective treatment packages, on the other hand, while implemented by the provider, may be less influenced by reimbursement (Shrank et al. 2005). Therefore, for the purposes of simulating the impact of alternative provider incentives, one might implement a ‘standard’ set of treatments that are not specific to any individual provider.

### 3.4 Claims Adjudication

In health care, the claims adjudication process refers to the process of determining the payer’s financial liability for its members’ medical claims. Our objective is not to replicate the nuances of an actual claims adjudication system in the real-world, but rather, to capture the features that are most important for estimating

the flow of money between payer and provider. To this end, we define a `ClaimsPayoutProcess`, belonging to a `Payer`, that contains payment processor objects (`PaymentProcessor`) and rule objects (`ClaimProcessingRules`) for how to sequence claims through the various payment processors. A *payment processor* is a core element of the claims adjudication process in our framework. It receives a claim from a provider, computes the eligible and reimbursable amount against that claim, and allocates any payments to the accountable provider(s).

A payer may possess multiple payment processors and each claim may be processed by different payment processors according to a set of rules. We currently provide templates for three types of payment processors: fee-for-service, episode or global capitation, and performance-based pay. Users of our framework may design their own variations of these conventional payment models by modifying various pre-defined, but general, parameters. Alternatively, they may design a completely new type of payment model by developing a custom implementation of the `IPaymentProcessor` interface. In Sections 3.4.1 and 3.4.2 we provide more details on the framework we use for specifying capitation and performance-based payments, respectively. We omit a description of fee-for-service payment processors since the logic behind them is relatively straightforward.

### 3.4.1 Capitation Payment

Capitation payments are any payments that provide a fixed reimbursement amount for a pre-specified category of care (e.g., primary care), or episode of care (e.g., hip replacement surgery), over a specified time period (Pham et al. 2010). The key attributes of a capitation payment processor are those that define the start and end of the episode. Typically, an episode is ‘triggered’ by certain medical procedure codes (e.g., Current Procedural Terminology (CPT) codes) appear in a claim. It may end after another procedure code appears in a claim, either immediately, or with some delay. An episode may only apply to certain patients and applicability may be determined by a patient’s diagnosis, as represented by the International Classification of Disease (ICD) codes that appear on the corresponding medical claim. Sometimes, to protect the provider against excessive losses, a *stop-loss* amount may be instituted, with the payer picking up some portion of the expenses beyond the stop-loss amount. Table 4 lists the core attributes for defining a capitation-based payment processors in our framework.

Table 4: Capitation Payment Processor Attributes.

Attribute	Description
Trigger Codes	The procedure(s) that trigger(s) the start of this procedure
End Codes	The procedure(s) that signal(s) the end of this episode
Duration	Fixed number of days
Duration Flag	If true, the duration above is applied when the end code(s) are encountered. Otherwise, it is applied at the onset of the episode
Eligible Providers	Participating providers
Covered Procedures	Set of procedures codes covered by this episode
Covered Diagnoses	Set of diagnosis codes covered by this episode
Covered Patients	Set of patients covered by this episode
Bundled Payment	Providers are paid this amount for the bundle of covered procedures
Stop-loss Value	Providers are insured against financial risks above the stop-loss value
Stop-loss Protection	The payer will pay this fraction of costs above the stop-loss value

In this table, the *duration* and *duration flag* are specified so as to allow the duration of an episode-based capitation to be defined relative to when the care was completed as opposed to started. For example, for hip replacement surgery, the bundled payment amount may apply to all care provided from the start of the initial surgical consultation until 30 days after the day of the surgery.

### 3.4.2 Performance-Based Payment

To implement a performance-based payment processor, our framework defines two key types of inputs, called *measures* and *metrics*. Performance measures are indicators of whether certain processes were performed, or outcomes were achieved, for an *individual* patient. Performance metrics are functions of one or more performance measures. While measures are determined at the level of the individual patient, a metric is a statistic based on one or more measures and collected over many patients. Metrics are evaluated over a specified time period and may be aggregated across, and attributed to, one or more providers. For both measures and metrics, we may categorize them as being *process* based or *outcomes* based.

To define a process measure in our framework, one creates an instance of a `PerformanceMeasure`, and must specify the relevant medical claim procedure code(s) as well as the criteria (e.g., patient diagnosis) for identifying the patients for whom the process should be performed. To define an outcome measure, one must specify the criteria for identifying eligible patients, as well as the conditions that defined a desired outcome. Examples of process and outcome measures are provided in Table 5. These examples are taken from UCare, the fourth largest health plan in the state of Minnesota (UCare 2010).

Table 5: Examples of Performance Measures.

Type	Description	Criteria
Process	Retinal eye exam performed	Members 18-75 years of age with diabetes (type 1 or 2) as identified through pharmacy and claims data.
Outcome	HbA1c in poor control	Same as above. In addition, patient's HbA1c must be > 9%.

To define a performance metric, one creates an instance of a `PerformanceMetric`. A simple example of a metric that is a function of the outcome measure in Table 5 would be the percentage of poorly controlled diabetic patients. In this case, the metric counts the number of patients who satisfy this measure and divides by a count of the total population of patients eligible for this metric. More generally, a metric may be defined so that a numerator count may only be incremented after a certain combination of measures are satisfied for an eligible patient. A different example of a metric is a cost-based metric that may be used in a shared-savings program. In this case, one could define a metric that sums the expenditures associated with all eligible process measures.

An *accountable unit* is the set of one or more providers contributing to the performance against a metric. In general, performance metrics may be a function of a provider's performance relative to other providers, or relative to past performance. Additionally, the payment rate may vary in a tiered fashion, with higher rates paid for higher relative performance. Table 6 lists the attributes of a performance metric. Finally, the core attributes of a performance-based payment processor are listed in Table 7.

Table 6: Performance Metric Attributes.

Attribute	Description
Accountable Unit	The provider(s) accountable for performance against this metric
Measures	The set of measures used to determine metric value
Value	The computed value of the metric

### 3.5 Integration of Evidence-Based Analytics

This section will describe how we envision our platform being used to integrate various data-driven mathematical methods to support the ongoing calibration of a simulation model against a payer or provider's medical claims and medical records databases. These analytics would serve to customize a simulation

Table 7: Performance-Based Payment Processor Attributes.

Attribute	Description
Start Date	The start date of the performance evaluation period
Expiration Date	The end date of the performance evaluation period
Accountable Units	The set of all eligible accountable units
Metrics	The set of performance metrics to be evaluated
Performance Tiers	An array of the performance thresholds that delineate changes in payment rate (or tier) for each metric
Payment Rates	An array of the payment rates to apply to each performance tier
Reference Values Map	A 1-to-1 map of metrics to reference values, where reference values are used to determine relative performance

model to the characteristics of a given population of patients and providers. There are four components of the simulation framework that are most likely to benefit from the use of data-driven models: the disease model, patient, medical intervention and provider components. The sources of data that could be used would include claims data, pharmacy data, lab results and patient health records. We recognize that data quality and availability will influence the type of analyses that can be performed effectively, which will also influence the degree to which the simulation model can be evidence-based in the sense of being strongly data-driven.

From the perspective of a Markov disease model, an approach called ‘state abstraction’ (Li et al. 2006) could be used to automatically extract the most relevant features for defining a compact set of disease states. While one might expect the set of disease states to be universal across patients, it may be the case that certain disease states are more relevant to certain patient populations. Additionally, the features defining the disease states may depend on the availability and characteristics of the data. Statistical methods (Welton and Ades 2005) could be used to estimate the transition probabilities between disease states, which may be dependent on various patient attributes. These attributes could correspond to various patient risk factors. Once the disease states have been defined, it is necessary to initialize patients as belonging to one of the identified disease states. Since patients may not cleanly fall into a unique state, rules may be needed to determine the most likely state of a patient.

While one could translate established clinical guidelines into the medical intervention structure described in §3.3, it may be of interest to identify patterns in the interventions being used in a particular provider population. Our hierarchical intervention structure is not unlike the structure found in a bill-of-materials for a product. Perlich and Rosset (2007) present a method for automatically identifying marketing product bundles using hierarchical clustering methods. Perhaps a similar approach could be applied to ‘bundling’ provider interventions. Finally, statistical estimation of the propensity of a provider to implement certain patterns of intervention could be useful for characterizing a provider’s behavioral model. The analysis might also be used to determine how such a propensity may be modified by both patient and provider attributes. Table 8 summarizes the examples just mentioned.

Table 8: Opportunities for Evidence-Based Analytics.

Framework Component(s)	Opportunity
Disease Model	Abstraction of relevant disease states
	Estimation of transition probabilities between disease states
Patients	Clustering and/or classification of patients into risk groups
	Classifying patients into a disease state, as defined by the disease model
Medical Interventions	Identifying treatment patterns/bundles
Provider	Estimation of propensity of provider to use identified treatment patterns

## 4 IMPLEMENTATION OF FRAMEWORK

The framework described in Section 3 provides a formal context for the various modeling components and establishes interfaces between them. In this section, we describe one implementation of our framework into a simulation flow using an existing modeling environment called The IBM Agent Building and Learning Environment (ABLE). We also describe how our simulation framework can be used to incorporate evidence in the form of claims, pharmacy, lab, and other medical records to create a ‘living’ simulation of the effect of alternative incentives in a health care delivery system.

### 4.1 Simulation Flow Sequence

In this section we describe an example simulation flow sequence. Please note that our framework is not limited to implementing only this type of sequence. Figure 3 illustrates an object (or agent) interaction diagram of a payment simulation flow. The simulation controller provides overall time sequencing and control of the simulation model and elements. These may include the primary actors in the simulation, `Patient`, `Provider`, and `Payer` and the supporting objects: `DiseaseModel`, `EncounterGenerator`, and `PaymentProcessor`. The disease model and payment processors were previously described in Sections 3.2 and 3.4, respectively. An *encounter generator* is an event processor that processes provider actions by ‘traversing’ a treatment package applied to a patient during a patient-provider encounter. It determines the next provider action that will take place, and then triggers the creation of any new claims from that encounter. Referring to Figure 3, the list below outlines the simulation flow.

1. The `SimulationController` directs the `DiseaseModel` to call its `progressDisease()` method, which considers the patient’s health state and treatment(s) and progresses the disease one time step (for example, a day, week, month, or year).
2. The `DiseaseModel` then updates the patient information so that the `Patient` reflects the current health state of the individual it represents.
3. The `SimulationController` directs the `EncounterGenerator` component to process encounters, or interactions, between patients and providers by calling its `generateEncounters()` method.
4. During an encounter, the `Provider` retrieves the contract from the `Payer`, examines the health state of the current `Patient` and prescribes a treatment using its `setPatientTreatments()` method.
5. Given the provider’s treatment decision, the `EncounterGenerator` generates a corresponding set of claims by calling its `generateClaims()` method and passes them to the `Payer`.
6. The `SimulationController` directs the payer to process any new claims.
7. The `Payer` adjudicates new claims using one or more payment processors.
8. Finally, the `Payer` transfers any reimbursement or other payments to the providers.

### 4.2 Implementation Using an Agent Building Software Platform

The IBM Agent Building and Learning Environment (ABLE) provides an agent-based modeling and simulation toolkit (Bigus et al. 2002) that can be used to easily represent our Healthcare incentives simulation framework. ABLE is a mature distributed multi-agent platform written in Java with a drag-and-drop GUI environment based in the Eclipse open-source development tool platform. Figure 4 shows a screen shot of the healthcare simulation model.

The mapping from the healthcare simulation framework to ABLE agents is straightforward. Each model element, such as a patient, provider, or payer, is represented by an agent. Similarly, elements such as interventions, encounter generators, and payment processors are also represented by Agents. These agents can be dragged-and-dropped onto the simulation model canvas and can be configured using associated

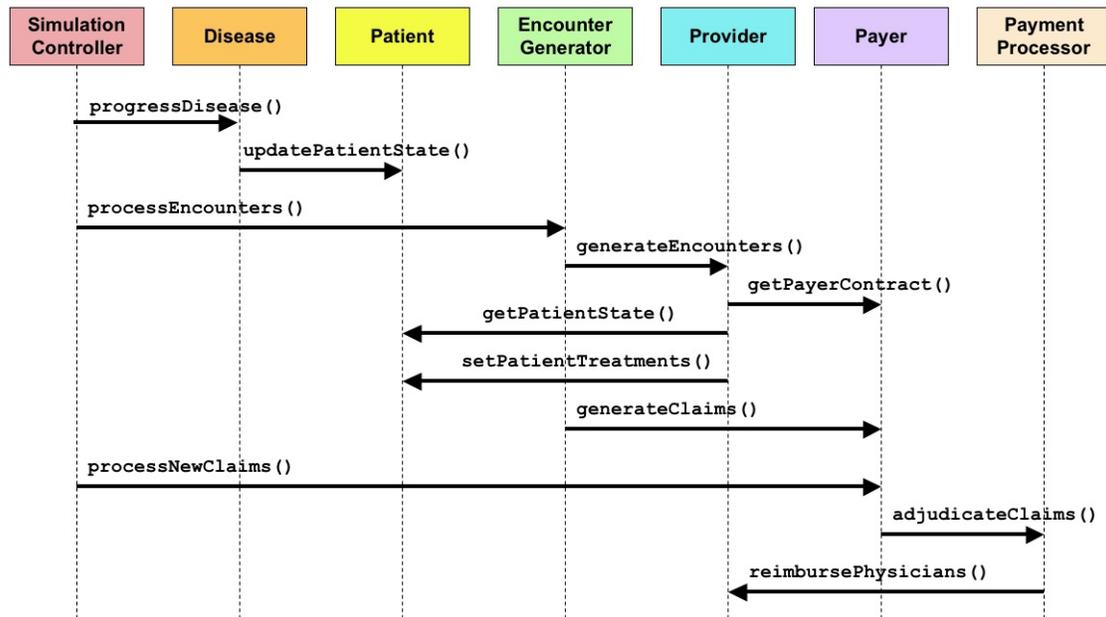


Figure 3: Example of Simulation Flow.

tabbed Properties sheets. Data can be loaded into agents from external files or databases using ABLE Import beans. This mechanism allows us to create agent populations of desired size with each agent having unique instance data read from the external file.

ABLE also provides a set of graphical Inspectors that allow text or charts to display internal state information during a simulation run. Simulation data can also be written out to external files or databases for post-hoc analysis using spreadsheets or business intelligence tools.

An important aspect of the ABLE simulation environment is that it is based on the ABLE distributed agent platform. This allows almost unlimited scalability, because additional agent pools (Java Virtual Machines) can be allocated to run selected populations of patients, payers and providers across multiple servers on a network. All of the inter-agent communication and messaging is handled by the agent platform services, so the simulation agents can focus on the behaviors of interest.

## 5 SUMMARY AND FUTURE WORK

We have presented a general framework for representing simulation model components for simulating a process of health care treatment, health progression, and health care related payments. We have implemented this framework on a Java-based software platform that provides a convenient, flexible and scalable modeling environment.

Planned extensions to our framework include incorporating the provider institution, patient behavior, global resources, and event-driven care processes. The current implementation of our framework only supports payment contracts between the individual provider and the payer. In reality, contracts may be written between payers and provider institutions (e.g., hospitals), and separate contracts written between provider institutions and the providers they employ. Adding this ‘layer’ of contracting increases the complexity of the incentives that influence provider behavior and requires one to model the behavior of the provider institution. We also plan to expand the role of the patient as an active decision-maker. This implies explicitly representing the behavioral models of the patient which we only treat implicitly by way of the ‘lifestyle’ attributes of a patient in the current framework. As an active decision maker, the patient can determine when to seek treatment, or comply to treatment, based on his health plan, resources and other attributes.

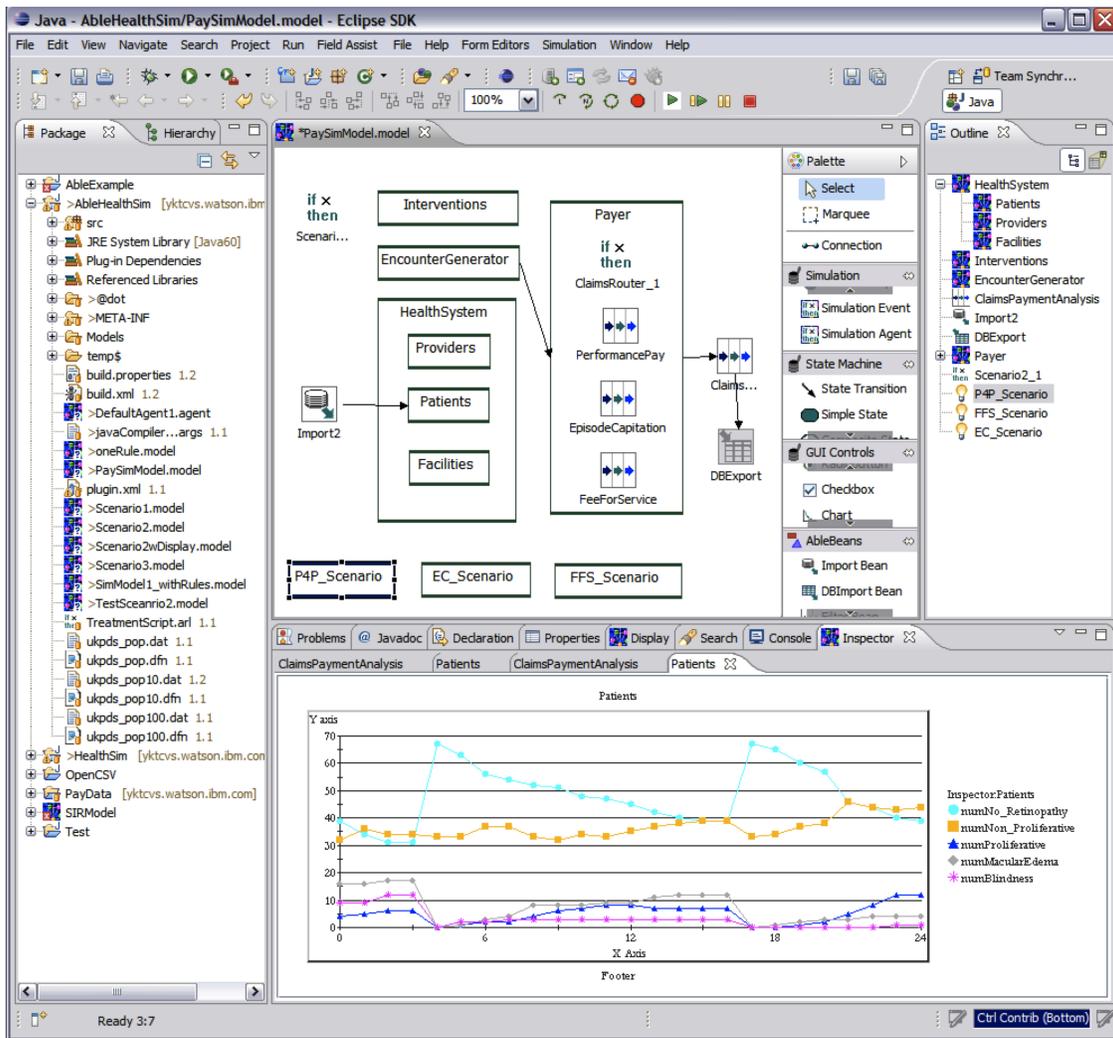


Figure 4: Screen Shot of Simulation Framework, as Implemented on ABLE Platform.

In our existing framework, any resource constraints influencing a decision-maker’s actions are ‘localized’ within the implementation of an agent’s behavioral model. By making resources global objects in the simulation, we will be able to model scenarios in which multiple agents share common, limited resources. While a global representation of resources might be a first consideration in most simulations that we see in operations research, our primary interest was in capturing clinical treatment selection based on rewards. Therefore, our initial focus has been on representing the reward mechanisms (e.g., the payment processors) in our framework. Even after global resource objects have been included in our framework, we do not anticipate using it to model very detailed dynamics of resource utilization (e.g., patient queues and schedules or medical supply inventory levels). Rather, we may use simpler, closed-form approximations of delay times, since this may be sufficient for estimating their impact on long term health outcomes.

The simulation flow described in §4.1 is currently implemented in batches by simulation year, and the care for one patient within each simulated year is delivered independently of the care for other patients. In many cases, this rather time insensitive analysis may be sufficient. However, in other cases, when the clinical care, diagnosis and disease progression is complex, with multiple parallel pathways (e.g., those that may be created by the patient referral process) and dependent clinical events, an event-based simulation

may be in order. We have yet to refine our framework and platform to implement these more complex processes.

Finally, in order to support evidence-based simulations, an important functionality to incorporate in our platform is that of providing a convenient method for owners of healthcare data (e.g., payers or providers) and owners of models for analyzing that data (e.g., researchers) to access each other's data/models. To do this, we will need to define a set of 'standard' simulation scenarios, and design generic data templates to support most models in this set.

## ACKNOWLEDGMENTS

The authors would like to acknowledge Sugato Bagchi and Jianying Hu for their contributions to the ideas and work presented in this paper.

## REFERENCES

- Aalen, O. O., V. T. Farewell, D. de Angelis, N. E. Day, and O. Nöel Gill. 1997. "A Markov model for HIV disease progression including the effect of HIV diagnosis and treatment: Application to AIDS prediction in England and Wales". *Statistics in Medicine* 16 (19): 2191–2210.
- Barhak, J., D. J. Isaman, W. Ye, and D. Lee. 2010. "Chronic disease modeling and simulation software". *Journal of Biomedical Informatics* 43 (5): 791–799.
- Bertakis, K. D., L. J. Helms, E. J. Callahan, R. Azari, and J. A. Robbins. 1995. "The Influence of Gender on Physician Practice Style". *Medical Care* 33 (4): 407–416.
- Bigus, J. P., D. A. Schlosnagle, J. R. Pilgrim, W. N. Mills III, and Y. Diao. 2002. "ABLE: A toolkit for building multiagent autonomic system". *IBM Systems Journal* 41 (3): 350–371.
- Cefkin, M., S. Glissmann, P. Haas, L. Jalali, P. Maglio, P. Selinger, and W. C. Tan. 2010. "Splash: A Progress Report on Building a Platform for a 360 Degree View of Health". 5th INFORMS Workshop on Data Mining and Health Informatics.
- Charfeddine, M., and B. Montreuil. 2010, December. "Integrated agent-oriented modeling and simulation of population and healthcare delivery network: Application to COPD chronic disease in a Canadian region". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hukan, and E. Yücesan, 2327–2339. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Choudhry, N. K., R. H. Fletcher, and S. B. Soumerai. 2005. "Systematic Review: The Relationship between Clinical Experience and Quality of Health Care". *Annals of Internal Medicine* 142 (4): 260–273.
- Christiansen, J., and A. Campbell. 2003. "HealthSim: An Agent-Based Model of Simulating Health Care Delivery". In *Proceedings of the International Microsimulation Conference on Population*.
- Dranove, D., and W. White. 1989. "Agency theory: new insights into the health care industry". *Journal of Medical Practice Management* 4 (3): 165–169.
- Fuloria, P., and S. Zenios. 2001. "Outcomes-adjusted reimbursement in a health-care delivery system". *Management Science* 47 (6): 735–751.
- Lee, D., and S. A. Zenios. "An Evidence-based Incentive System for Medicare's End-Stage Renal Disease Program". *Management Science*. Forthcoming.
- Li, L., T. J. Walsh, and M. L. Littman. 2006. "Towards a Unified Theory of State Abstraction for MDPs". In *Ninth International Symposium on Artificial Intelligence and Mathematics*, 531–539.
- McGuire, T. 2000. *Physician agency*, Volume 1 of *Handbook of Health Economics*, Chapter 9, 461–536. Elsevier.
- Montreuil, B., M. Charfeddine, O. Labarthe, and A. C. André.. 2007. "A generic agent-based framework for simulation in distributed healthcare systems". International Conference on Industrial Engineering and Systems Management.

- Perlich, C., and S. Rosset. 2007. "Identifying Bundles of Product Options using Mutual Information Clustering". In *SDM*. SIAM Conference on Data Mining.
- Pham, H. H., P. B. Ginsberg, T. K. Lake, and M. M. Maxfield. 2010, January. "Episode-Based Payments: Charting a Course for Health Care Payment Reform". Technical Report 1, National Institute for Health Care Reform.
- Shrank, W., S. L. Ettner, P. H. Slavin, and H. J. Kaplan. 2005. "Effect of Physician Reimbursement Methodology on the Rate and Cost of Cataract Surgery". *Archives of Ophthalmology* 123 (12): 1733–1738.
- Silverstein, M. D., E. V. Loftus, Jr., W. J. Sandborn, W. J. Tremaine, B. G. Feagan, P. J. Nietert, W. Harmsen, and A. R. Zinsmeister. 1999. "Clinical course and costs of care for Crohn's disease: Markov model analysis of a population-based cohort." *Gastroenterology* 117 (1): 49 – 57.
- Minnesota UCare 2010. "UCare 2010 Pay for Performance Program Measures and Anticipated Payment Rates".
- Welton, N. J., and A. E. Ades. November/December 2005. "Estimation of Markov Chain Transition Probabilities and Rates from Fully and Partially Observed Data: Uncertainty Propagation, Evidence Synthesis, and Model Calibration". *Medical Decision Making* 25 (6): 633–645.
- Yaesoubi, R., and S. Roberts. 2011. "Payment contracts in a preventive health care system: A perspective from Operations Management". *Journal of Health Economics*. Forthcoming.

#### **AUTHOR BIOGRAPHIES**

**JOSEPH P. BIGUS** is a Senior Technical Staff Member at the Thomas J. Watson Research Center, where he is the project leader on the ABLE research project. He is a member of the IBM Academy of Technology and is an IBM Master Inventor, with over 25 U.S. patents. He was an architect of the IBM Neural Network Utility and Intelligent Miner for Data products. Dr. Bigus received his M.S. and Ph.D. degrees in computer science from Lehigh University, his MBA from U. Mass., Amherst, and a B.S. in computer science from Villanova University. His current research interests include agent-based modeling and simulation and their application to understanding healthcare systems payments and incentives. His e-mail is [bigus@us.ibm.com](mailto:bigus@us.ibm.com).

**CHING-HUA CHEN-RITZO** is a Research Staff Member in the Business Analytics and Mathematical Sciences Department at the IBM T.J. Watson Research Center. She received a B.S. degree in physics from the University of North Carolina at Chapel Hill in 1997, an M.S. degree in Architectural Engineering in 1999, and a dual-title Ph.D. degree in Business Administration and Operations Research in 2006, both from Penn State University. She has been at IBM since 2005, where she has worked on applied mathematical optimization and simulation for applications in services and manufacturing operations. Her e-mail address is [chenritzo@us.ibm.com](mailto:chenritzo@us.ibm.com).

**ROBERT SORRENTINO** is a Research Staff Member in the Healthcare Transformation Department at the Thomas J. Watson Research Center. He received a S.B. degree in physics from Massachusetts Institute of Technology in 1973, and M.D. degree from Stony Brook University School of Medicine in 1978. He practiced Emergency Medicine for fifteen years at a high-acuity, inner-city hospital in Arizona. In addition to a long-term role as an independent healthcare management consultant for Mercer Health and Benefits, Dr. Sorrentino was Chief Medical Information Officer at CareMore Health Plan, and its MSO affiliate, CareMore Medical Enterprises, from 2002 through its acquisition by JPMorgan Partners in 2006. He subsequently served as Chief Medical Officer for ARTA Medicare Healthplan, and its affiliate MSO, Western Medical Management in Southern California. Dr. Sorrentino later joined IBM at the Thomas J. Watson Research Center, where he has worked on advanced healthcare analytics, and innovative payment models for healthcare. He is an author or coauthor of 6 patents and 4 technical papers. His e-mail address is [sorrentino@us.ibm.com](mailto:sorrentino@us.ibm.com).