

Chapter 2

Personalized Medicine: Integrating Individual Exposure and Response Information at the Bedside

Diane R. Mould and Lawrence J. Lesko

2.1 Introduction

Historically, therapeutic agents were dosed using the same dose for all patients (“flat dosing”), sometimes dosed on body size (weight or body surface area, BSA), or adjusted based on key patient factors (covariates), such as degree of renal impairment. For some agents such as warfarin, the risks associated with both over and underdosing are substantial, and genetic markers can be used to refine the starting dose and the dose increments in order to safely achieve the international normalized ratio (INR) target range and subsequent clinical effect. However, owing to differences in tolerance such as with antineoplastic agents, or effects, such as with antihypertensive agents, adaptive dosing where doses are adjusted based on observed response (“adaptive dosing”) is also used. Individualizing drug therapy, or tailoring the selection of both the drug and the dose for a specific patient, has been a long-held objective of physicians and other health-care providers. As stated in a recent review of the history of individualized medicine (Lesko and Schmidt 2012), “personalized medicine is an evolution, not a revolution.”

Personalized medicine is expected to optimize the benefit and minimize the harm of medical interventions on a patient-by-patient basis. Thus, the goal of personalized medicine is to identify patient characteristics predictive of response to therapy and to use this information to provide a therapeutically optimal dose for each patient, or patient subgroups, based on their individual characteristics (Conti et al. 2010). Examples of patient characteristics that may affect drug exposure and response, and subsequently require individualization of treatment and dose include age, body weight, race, sex, organ function (e.g., hepatic and renal function), and various types of biomarkers, such as biochemical, disease markers, and genomic

D. R. Mould (✉)

Projections Research, Phoenixville, PA, USA

e-mail: drmould@pri-home.net

L. J. Lesko

Center for Pharmacometrics and Systems Pharmacology, University of Florida, Orlando, FL, USA

© American Association of Pharmaceutical Scientists 2014

S. Schmidt, H. Derendorf (eds.), *Applied Pharmacometrics*, AAPS Advances in the Pharmaceutical Sciences Series 14, DOI 10.1007/978-1-4939-1304-6_2

markers. The goal of personalized medicine aligns well with that of population pharmacokinetic and pharmacodynamic (PK/PD) modeling, which includes identification of covariate factors that are predictive of heterogeneity and uncertainty in drug exposure and/or response.

The utilization of biomarkers for patient care decisions has been limited by the lack of decision-support tools for practitioners to facilitate integration of biomarker data with other patient specific information to generate a treatment recommendation (Zineh and Huang 2011). PK/PD modeling enables integration of multiple patient characteristics in a drug-specific framework, and recently has been combined with web-based applications that provide a user-friendly interface, or “dashboard” for including patient-specific inputs, updating the associated models, and summarizing and visualizing the data and model-based dose predictions (Barrett et al. 2008). Such dashboard systems have the potential to offer an improved and convenient means for health-care provider to tailor treatment for an individual patient, particularly for drugs with high variability in exposure or a narrow therapeutic window.

2.1.1 Current Dosing Paradigms

There are numerous approaches to developing dose regimens for therapeutic agents, but the most common are the “flat” dose (e.g., all patients receive the same dose), with dosing based on body size also being a common dose metric. In addition, dosing is often stratified based on covariates, such as genotype or organ function. Adaptive dosing, where doses are increased or decreased based on observed effect is also used.

One of the issues with the flat dose option is that the exposure and/or response to a given dose is often highly variable. This variability can arise from differences in the PK, such as in genetic subpopulations, that rapidly clear a drug, or can be due to differences in the PD related to a given plasma drug concentration. If there are factors, such as an effect of weight on clearance, then small patients will tend to be overdosed and patients with high weight will be underdosed using this dose adjustment strategy. Figure 2.1 depicts two hypothetical patients’ response to the same dose of a drug. Depending on the therapeutically desired response, these patients may need to receive higher (in the case of the patient with a weaker response) or lower (in the case of the patient with a greater response) doses.

Dosing based on body size is a common approach. However dosing on a mg/kg basis often results in subtherapeutic exposures in low weight patients, particularly pediatric patients (Anderson and Holford 2013; Xu et al. 2013) because the relationship between drug clearance and weight (if it exists) is rarely linear due to the differences in the ratio of clearance organ size to overall body weight. This finding has been confirmed for many compounds, including infliximab (Xu et al. 2012; Fasanmade et al. 2011). The US Food and Drug Administration (FDA) has written a guidance document for industry on dose selection for the minimum recommended dose for first time in humans (FTIH) studies (Guidance 2014) that suggests selection of initial doses based on body weight (e.g., mg/kg) in order to scale exposure observed in nonclinical studies to safe levels in humans. While this document is




Fig. 2.1 Examples in differences between patients in response to a specific dose. In this figure, two hypothetical patients were administered the same dose of drug with the goal of lowering the measured response indicator. The patient response represented by the *black line* is a patient with a strong response while the patient whose response is represented by the *red line* has a weak response. Thus the latter patient may need a higher dose or a different treatment for their disease

not specifically aimed at providing guidance for dose selection, many marketed monoclonal antibodies (mAbs) are labeled for dosing on a mg/kg basis. A recent review (Mould and Green 2010) found that only three of 26 marketed mAbs had a clearance that was linearly related to weight, eight were dosed on a mg/kg basis and two of these had no weight effect identified on the clearance. Doses based on BSA are similarly problematic. Egorin published a review on BSA-based dosing for antineoplastic agents (Egorin 2003). The variability in exposure with this dosing approach is not always improved as compared to “flat” dosing.

A stratified dose approach where flat doses are administered over specified ranges of body weight, or over and under a given mg/kg weight, is often the best way to ensure appropriate dosing when body size impacts clearance, and may be particularly relevant for pediatric patients (Xu et al. 2013). This approach has the benefit of reducing the overdosing and underdosing seen with flat dosing and dosing based on body size, either weight or BSA.

Some compounds such as epoetin (a biologic agent used to treat anemia) are dosed based on specific hemoglobin measurement. The dose algorithm is complex however, and although the approach works well to control hemoglobin, the complexity of the dose strategy can give rise to dose errors, and dose adjustments takes time to determine. Computer-based dose support has been shown to improve the percentage of patients staying within the target range of hemoglobin, often with a lower dose than the manual adjustment provided (Ho et al. 2010), and with substantially increased staff efficiency without having a negative impact on safety

(Miskulin et al. 2009). Thus, computer-guided dosing may have a substantial impact on optimizing patient management of their therapies.

2.1.2 Definition of a Dashboard

A dashboard is a user interface that, like a dashboard in a car, organizes and presents information so that is easy and quick to read and interpret. Software packages that integrate information from multiple components into a unified display are referred to as dashboards. For example, patient management dashboards might obtain information from electronic medical records, laboratories, and through clinician and patient input and present it as though it all came from the same source. Hewlett Packard (HP) developed the first dashboard system, which began as a tool for customizing Windows desktops. Called “Dashboard,” the HP product was later acquired by Borland and then a company called Starfish (Dashboard 2014).

“Dose calculators” have been in existence since the late 1950s, although the majority of these early systems were to calculate radiological doses (Sivyer 1959). Until recently, the computational needs of individualized dosing were limited, although the application of Bayesian forecasting has been shown to result in therapeutic improvements. For example, application of Bayesian-based dosing substantially increased the number of patients whose trough phenytoin levels were within the target range (63.6% of the phenytoin troughs from the Bayesian forecasting group, compared with 34.0% in the conventional dose adjustment group) (Tobler and Mühlebach 2013). One of the earlier dashboard systems in clinical use focused on antineoplastic dosing for pediatrics (Barrett et al. 2008), and the number of dashboard systems has grown over time.

A related topic that will not be covered in detail here is the emergence of computerized clinical decision support systems (CDSS). Papier (2012) defined these as “an interactive system allowing input of patient-specific information and providing customized medical knowledge-based results via automated reasoning, for example, via a set of rules and/or an underlying logic, and associations.” These systems generally are not based on an underlying population model but embody collected clinical expertise which is compared to a patient’s symptoms using methods such as rule-based or fuzzy logic algorithms (Domínguez Hernández et al. 2013). Like dashboard systems, they are a growing area of research seeking to maximize the use of prior knowledge for an individual patient.

2.1.3 Relationship to Population Models

Dashboard systems are generally built around a population model (Mould and Upton 2012). The population model is essentially an embodiment of the current state of knowledge about the PK or PD of a drug and generally includes three key components:

1. The structural (base) model (e.g., a one-compartment PK model) that provides a (ideally) mechanistic description of the time course of a measured response.

2. Stochastic (probability) models that describe the distribution of unexplained variability in the observed population, such as between-subject variability (BSV) or residual variability (RUV).
3. Covariate models that quantitate the influence of explainable factors such as demographics or disease on individual time course of the response.

The dashboard system is intended to be a user-friendly system for accessing the model to forecast exposure or response for an individual patient, with the predictions of the model refined by incorporating information about each individual patient into the model. The greater the size and diversity of the database used to construct the underlying model, the greater the chance that the model will be able to return useful and accurate individual predictions for most patients. However, it should be recognized that there are limitations to forecasting using population models. The underlying assumption is always that a particular population model will continue to describe data from a patient into the future, and that the model captures all important sources of variability (both explainable and unexplainable). When the underlying assumptions do not apply to a given patient, the model predictions have the potential to be in substantial error. For example, a given model of drug PK may have been appropriate for a patient in the past, but if the patient has a cardiac infarct (with major reductions in cardiac output with a subsequent impact on drug clearance), the forecast concentrations from the model may be substantially underestimated.

2.2 Individualized Forecasts

There are two main mechanisms by which individual data can be used to refine the predictions of a population for a particular patient. These are via covariate relationships identified during the model building process, and by Bayes updating on model parameters based on individual data. Figure 2.2 shows an example of how both methods can work together to improve the forecast for an individual patient.

2.2.1 *Covariate Effects*

Covariates (e.g., age, sex, renal function) modify the value of a model parameter (e.g., clearance, CL) depending on the value of the covariate (e.g., body weight on CL in Fig. 2.2). The inclusion of a covariate relationship in a model will generally imply that the model provides a better description of the data and that the unexplained BSV of the associated parameter is reduced. Thus, covariate factors effectively convert unexplainable variability to explainable variability at the population level, and reduce the uncertainty in the values of the model parameters for individual patients for whom covariate values are known (Fig. 2.2). However, depending on the drug and the dataset used to develop the model, the contribution of covariates to reductions in unexplainable variability can vary from nothing (i.e., no covariates identified) to modest or substantial contributions. When no covariates are found, this implies that

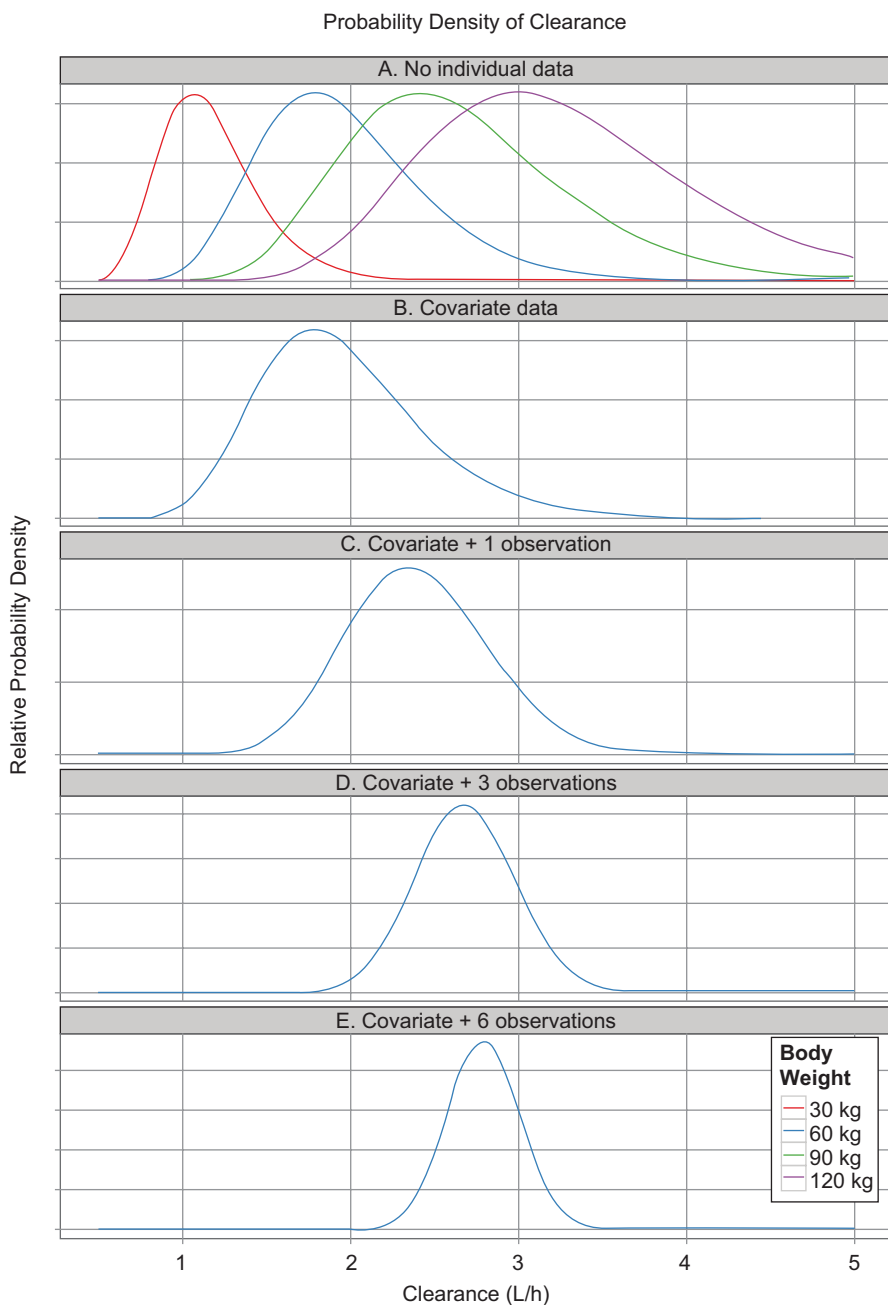


Fig. 2.2 An example of the contribution of individual data to Bayes forecasts. Below is an example model to determine steady-state drug concentration (C_{ss}) of a chronically administered drug: $C_{ss} = \text{DoseRate} / (\text{CL} * (\text{WT}/70)^{0.75})$. CL is a log-normally distributed population parameter with a population value of 2 L/h and BSV of 25%. Patient body weight (WT) is a covariate affecting

the factors causing variability between patients have either not been identified, were not possible to measure in sufficient numbers of individuals, or were not available at all in the analysis dataset. It is important to note that the ability to identify predictive covariates is dependent on both the method used to evaluate the data (Wählby U, Jonsson EN, Karlsson MO. Assessment of actual significance levels for covariate effects in NONMEM. *J Pharmacokinet Pharmacodyn.* 2001; 28(3):231-52) and the approaches used during modeling (Mould and Upton 2013). However, even when covariates are identified, many agents still have considerable unexplained variability, which limits the use of patient covariates in individualizing dosing.

2.2.2 *Bayes Update of Models with Individual Data*

Most of the currently available software packages for individualizing therapy utilize Bayesian methods to help predict future response to a given dose regimen. In general, such software packages use a mixture of Bayesian updating, Bayesian forecasting, and Bayesian model averaging. Bayesian inference is a method in which Bayes' rule is used to update the probability estimate for a hypothesis as additional data are obtained. Bayesian updating is particularly important in the dynamic analysis of data collected sequentially over time.

Bayesian updating uses a model that not only describes the time course of exposure and response but also includes terms describing the unexplained (random) variability of exposure and response. It involves applying a “prior” (which is called a prior because it reflects the underlying information derived from previous evaluations) to form the underlying hypothesis. The prior distribution is the distribution of the parameter(s) before any new data are observed and is usually developed in a separate analysis. The prior therefore is the series of mathematical models describing exposure and response following administration of a drug. The sampling distribution is the distribution of the observed data conditional on its parameters. This is also termed the likelihood, especially when viewed as a function of the parameter(s). The marginal likelihood (also called a “posterior”) is the distribution

CL via an allometric relationship, where the standard body weight is 70 kg. DoseRate is the average steady-state dose rate—set at 10 mg/h, proportional residual error for the model was 20%. We wish to forecast the clearance of the drug (so that individual C_{ss} can be estimated). *Panel A* shows the probability densities for CL for the case where no individual patient information is known (densities are normalized to the same peak value for clarity). There are a variety of possible distributions for CL, depending on the unknown body weight of the patient. *Panel B* shows the distribution of CL with covariate data. The patient has a weight of 60 kg, eliminating other candidate distribution curves. *Panel C* shows the distribution of CL with covariate data and a single observation of C_{ss} in the patient. In this case, C_{ss} was found to be 4 mg/L, which was lower than the expected value of 5.6 mg/L for a 60-kg subject. The distribution therefore moves to the right, reflecting higher individual clearance and becomes narrower, reflecting more certainty about the individual patient value of CL. *Panels D* and *E* show the distribution of CL with covariate data, and 3 and 6 observations of C_{ss} in the patient, respectively. Note that as more individual data are available, the uncertainty in the distribution of CL reduces (i.e., the distributions are narrower) via Bayesian learning. Adapted from Mould DR, Upton R, Wojciechowski J. Dashboard Systems: Implementing Pharmacometrics from Bench to Bedside. *AAPS J* ePub June 2014 with permission

Table 2.1 The balance between the prior and the data

Factors that favor the prior	Factors that favor the data
Few data points	Many data points
High residual error	Low residual error
Low population variability	High population variability

of the observed data marginalized over the parameter(s). Thus, Bayes’ rule can be applied iteratively. That is, after observing data, the resulting posterior probability can then be treated as a prior probability, and a new posterior probability computed from the next set of new evidence. This procedure is termed Bayesian updating or sometimes “Bayesian learning” (Gill 2008).

However, rather than estimate the parameters for the model based solely on the patients data, Bayes’ theorem is implemented to balance the contribution of new data and prior knowledge in the estimation of the model parameters for the individual (see Table. 2.1). Thus, a single data point in an individual is given less weight in the fitting process if it deviates substantially from what has happened before, but is given more weight as additional data points support the finding. Similarly, a parameter value is given less weight in the fitting process if it deviates substantially from the prior values inherent in the population model (see Fig. 2.2). From a Bayes perspective, the interpretation of a data point is seen to have contributions respectively from the truth (the underlying process, described by a model), the errors (intraindividual, interindividual, interstudy, residual, etc.), and the prior knowledge:

$$\text{data} = \text{truth} + \text{error} + \text{prior knowledge.}$$

The updating process involves sampling parameters from the prior distribution and calculating the expected response based on the model, then comparing the difference between the model expectation and the observed data. This difference is referred to as the objective function. The parameters are then adjusted based on the objective function and the new parameters are tested. This process runs iteratively until the objective function is as low as possible (referred to as “minimizing the objective function”) suggesting that the parameters are the best to describe the current data. The result of Bayesian updating is a set of parameters conditional to the observed data balanced by the application of the principles of Bayes’ theorem.

Bayesian model averaging (Hoeting et al. 1999) offers a systematic method for checking the robustness of one’s results to alternative models. The standard practice of selecting a single model from some class of models, and then making inferences based on this model ignores model uncertainty, can impair predictive performance and overestimate the strength of evidence for predicting dose–exposure relationships. Bayesian model averaging allows model uncertainty to be incorporated into inference. The basic idea behind Bayesian model averaging is to make inferences based on a weighted average over model space which includes several models. This approach accounts for model uncertainty in both predictions and parameter estimates. The resulting estimates incorporate model uncertainty and thus may better reflect the true uncertainty in the estimates.

Table 2.2 Overview of selected dashboard systems

Software	Bayesian updating	Bayesian forecasting	Bayesian averaging	Dose	Drugs	Website
Abbottbase (Wong et al. 2013)	Yes	Yes	No	Yes, to an AUC	Aminoglycosides	NA
Drugcalc (García et al. 1994)	Yes	No	No	Yes to an AUC	Aminoglycosides	http://www.testandcalc.com/drugcalc/index.asp
Dosecalc (Mohan et al. 2004)	No	No	No	Yes to an AUC		NA
MW/Pharm (Usman et al. 2013)	Yes	Yes	No	Yes	180 drugs	http://www.mwpharm.nl/main.htm
CHOP Pediatric Knowledgebase Dashboard (Barrett et al. 2008)	Yes	Yes	No	Yes through forecast	Pediatric oncology (methotrexate)	http://pkb.chop.edu/index.php
NZ FirstDose Dashboard (Holford et al.)	No	Yes	No	Yes	Amikacin and vancomycin	http://www.firstdose.org/
TCIworks (Wong et al. 2013)	Yes	No	No	No	Gentamycin and enoxaparin	http://www.tciworks.info/
Warfarin dosing	No	No	No	Yes based on covariates	Warfarin	http://www.warfarin-dosing.org/Source/Home.aspx
Baysient dose evaluation system (Mould et al. 2013)	Yes	Yes	Yes	Yes, multiple	Any	http://www.baysient.com

NA not applicable

Bayesian forecasting (Elliott et al. 2006) then involves using the updated individual parameters to forecast the likely exposure and response that a given patient will exhibit with varying dose regimens based on the individual parameter estimates obtained via Bayesian model averaging and Bayesian updating. However, when the software does not have the capacity to do Bayesian updating, then the forecasting is generally based on the patient covariates which is generally less precise.

The majority of dashboard systems available are for use with aminoglycoside antibiotics and warfarin although there is one (Knowledgebase) that deals with dosing pediatric oncology. A list of several currently available systems is provided in Table 2.2. As can be seen, these systems utilize varying aspects of Bayesian methods to determine an individualized dose.

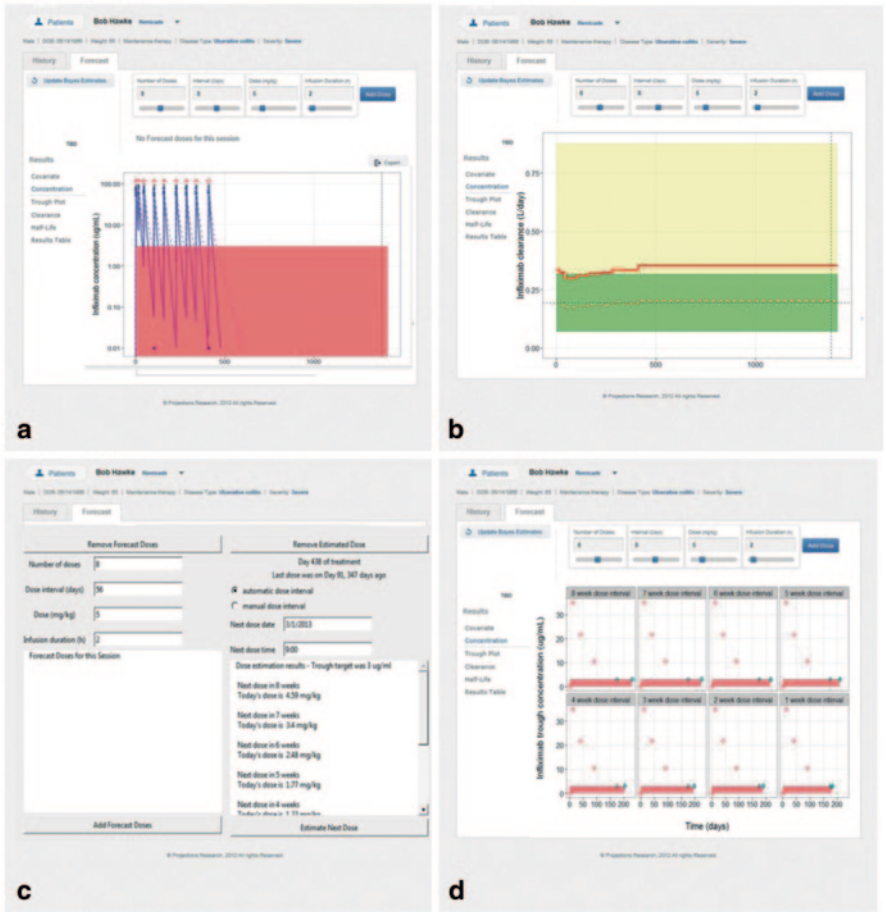


Fig. 2.3 Common dashboard screens. *Panel A* shows the agreement between the model with Bayesian updated parameters (blue lines) and the observed data (blue dots). Note that the concentrations are within the red shaded region suggesting that this patient is not at or above the target level. *Panel B* shows the patients individually estimated clearance over time. The green region is ± 3 standard deviations of a typical patient with those covariates. The fact that this patient's clearance is in the yellow shaded area suggests that the doses and frequency needed to maintain this patient at the target level will probably exceed the labeled recommendations. *Panel C* shows the dose optimization screen where the clinician can determine either an appropriate regimen (e.g., dose and interval), or can enter the next patient visit to ensure the patient is adequately covered during the interval. *Panel D* shows the expected troughs from all of the recommended dose regimens to ensure they are high enough to achieve a desirable response (<https://www.baysient.com>)

2.3 Dashboard Systems

In general, dashboard systems have several components including: (1) patient data management, (2) updating/forecasting, and (3) dose recommendations. Figure 2.3 shows screen shots of the results from Bayesian updating and forecasting as well

Applied Pharmacometrics

Schmidt, S.; Derendorf, H. (Eds.)

2014, XIII, 565 p. 161 illus., 94 illus. in color., Hardcover

ISBN: 978-1-4939-1303-9