We do *not* need to use the correct values of k_3 and k_4 to estimate $\boldsymbol{\phi}$; even if we use values of k_3 and k_4 that are off by an order of magnitude from the values used to generate the data, the estimate of $\boldsymbol{\phi}$ is the same as in Equation 9.66. In fact, we must design a new experiment to determine the values k_3 and k_4 if they are of interest, because they cannot be found using data shown in Figures 9.23–9.25. Please note that these particular conclusions depend on the particular values of the rate constants we used to generate the data. If very different parameter values are chosen, the method of analysis can be used again but the model reduction may change. □

The complexity of this model probably precludes us from guessing quickly and intuitively that parameter pair k_3 and k_4 (or k_3 and k_1) cannot be determined from the experimental data, and we should estimate certain ratios and differences of rate constants instead of the rate constants themselves. By estimating all parameters, and analyzing the eigenvalues and eigenvectors of the Hessian, we have a systematic approach to develop such understanding fairly quickly when presented with a new problem of interest.

9.3 An Industrial Case Study

Reactor modeling in the industrial environment is a challenging task. When successful models are constructed, however, the payoff can be large, resulting in more efficient reactor operation, more consistent and higher quality product, and improved reactor designs. Given the difficulty of the modeling challenge in the best industrial circumstances, the least we should expect is that the computational and numerical procedures being employed are efficient and reliable, and are not compounding the difficulty of extracting models from the available data. To conclude this chapter we examine a small prototypical case study that illustrates some of the challenges we can expect when we model data from industrial facilities. This study was conducted in collaboration with colleagues at Kodak and more details are available [26].

End-point problems. In many chemical reactions, two main reactants are combined to yield a primary desired product. A small quantity of one of the reactants often remains at the end of the reaction due to batch-to-batch variability in the purity and reactivity of the starting materials, and variability in the rate of side reactions. In some cases, an excess of one reactant at the final time does not present a problem and the opportunity for improvement with better control is proportional to

the quantity remaining, which is usually small. In these cases, an excess of one reactant is commonly added to ensure complete consumption of the other reactant.

In some cases, however, a small amount of neither starting material is tolerable. These materials may be difficult (impossible for practical purposes) to separate from the products, and their presence even in small quantities may prevent further processing steps from occurring. The excess addition of one reactant could also result in undesired reactions, usually after the limiting reagent is exhausted. In these cases, the cost of having the unreacted materials present in some amount greater than a small threshold is the cost of the entire batch of chemicals. This general problem in which neither reagent can exist at the end of the reaction is called the "end-point control" problem.

The reaction of interest is the dehalogenation of a dihalogenated starting material to form the divinyl product, which is used in photographic film production. It is assumed that the halide groups are removed from the starting material in two consecutive reactions:

$$A + B \xrightarrow{k_1} C + B\cdot HX$$

$$C + B \xrightarrow{k_2} D + B\cdot HX$$

A	$XH_2CCH_2RCH_2CH_2X$
B	Organic base
C	$H_2C{=}CHRCH_2CH_2X$
D	$H_2C{=}CHRCH{=}CH_2$

The dihalogenated starting material (A) loses HX to the base (B) to form the mono-halogenated intermediate (C), which subsequently loses HX to the base to produce the desired final product (D). Over-addition of base causes polymerization and loss of the batch.

Modeling. The reaction is carried out in a well-stirred semi-batch reactor as shown in Figure 9.29. The reactor is initially charged with a weighed amount of component A and component B is added. The ma-

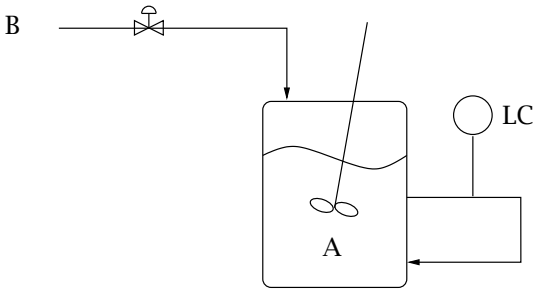

Figure 9.29: Semi-batch reactor addition of component B to starting material A.

terial balances for the semi-batch reactor are

$$\begin{aligned}
\frac{dV_R}{dt} &= Q_f(t) \\
\frac{d(c_A V_R)}{dt} &= -k_1 c_A c_B V_R \\
\frac{d(c_B V_R)}{dt} &= Q_f c_{Bf} - (k_1 c_A c_B + k_2 c_C c_B) V_R \\
\frac{d(c_C V_R)}{dt} &= (k_1 c_A c_B - k_2 c_C c_B) V_R \\
\frac{d(c_D V_R)}{dt} &= k_2 c_B V_R
\end{aligned} \tag{9.67}$$

in which Q_f is the volumetric flowrate of base and c_{Bf} is the feed concentration of B. The primary assumptions in this model are: isothermal operation, negligible volume change upon reaction or liquid chromatograph (LC) sampling, perfect mixing, negligible side reactions such as polymerization of the C and D, and reaction with impurities and inhibitors in the starting material. The initial conditions for the ODEs are:

$$\begin{aligned}
V_R(0) &= V_{R0} \\
n_A(0) &= n_{A0} \\
n_B(0) = n_C(0) &= n_D(0) = 0
\end{aligned}$$

It is assumed that c_{Bf} and V_{R0} are known precisely, and therefore the unknown model parameters are given by $\boldsymbol{\theta} = [k_1 \; k_2 \; n_{A0}]^T$. Although the initial A charged to the reactor is weighed, the parameter n_{A0} is chosen to be adjustable to account for the unknown level of impurities

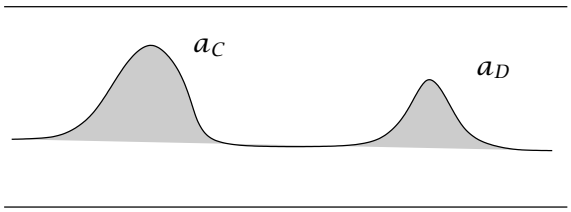

Figure 9.30: Depiction of an LC curve for determining the concentration of intermediate C and product D.

and the neglected side reactions. Notice that making good decisions of this type depend on experience and judgment, and reflect more the art rather than the science of reactor modeling.

The feed flowrate of base is measured by a mass flow meter. For composition analysis there is an on-line LC that draws a sample every 8 to 10 minutes and, after a 7-minute delay, reports the relative amounts of C and D. The unknown parameters, $\boldsymbol{\theta}$, are estimated with the dynamic LC data. The LC-detector wavelength is set to detect the R-vinyl bond. Therefore only the C and D species show peaks in the LC output as depicted in Figure 9.30. The areas of the two peaks are related to the molar concentrations by

$$a_D = 2k_{\text{lc}}c_D \tag{9.68}$$

$$a_C = k_{\text{lc}}c_C \tag{9.69}$$

in which k_{lc} is the proportionality constant for the LC. By calculating normalized areas this constant can be removed,

$$y(t) = \frac{a_C}{a_C + a_D} = \frac{c_C}{c_C + 2c_D} \tag{9.70}$$

This normalized C peak area, y, is used for the parameter estimation. The parameter estimation is performed by minimizing the following relative least-squares objective function,

$$\Phi = \sum_i \left(\frac{\tilde{y}_i - y_i}{y_i} \right)^2 \tag{9.71}$$

in which $\tilde{y}_i$ is the measured normalized C peak and y_i is the model prediction at the ith sampling time.

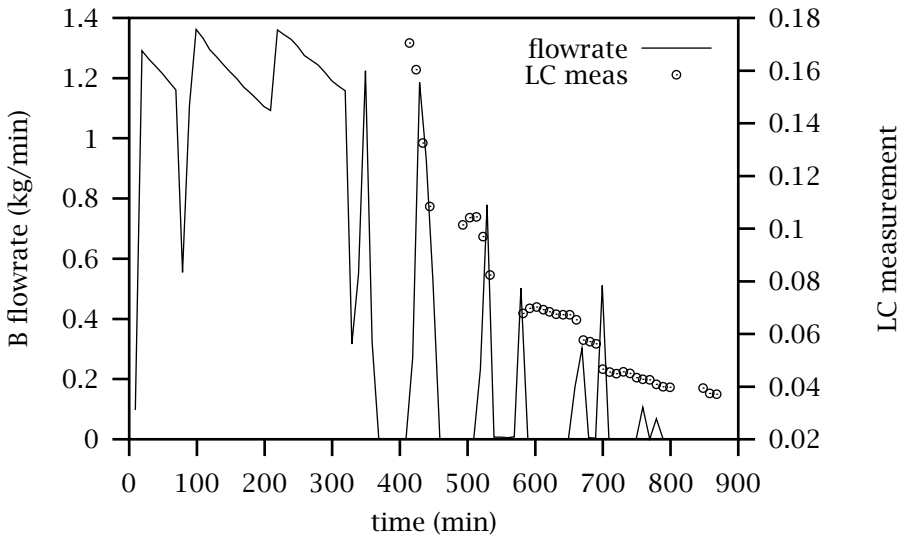

Figure 9.31: Base addition rate and LC measurement versus time.

Conventional manufacturing procedure. Normal operating procedure is to charge the reactor with some (only approximately known) amount of A, and add an initial amount of B that is sure not to overshoot the end point. An example of a typical B addition profile is given in Figure 9.31. After the initial B is added, the LC is switched on and the operator waits until the readings stabilize. The operator then sees how much C is remaining, and — based on experience — adds more B. The objective is to add enough B to consume all but 3% of the A. Ideally the operator would like to consume all the A, but the target is set at 3% to allow a margin for error. The penalty for overshoot is so high that only conservative addition steps are generally ever taken. After making the addition, the LC readings are again allowed to stabilize, and the operator again checks to see how close he is to the target before making another addition. This cycle repeats until the operator is satisfied that he is close enough (between 2–4%) to the target. As can be seen in Figure 9.31, the operator required 7 additions and about 500 minutes after first turning on the LC before he had determined how much B was required to reach end point.

Parameter estimation. If we perform parameter estimation with these data we achieve the model fit shown in Figure 9.32 The model predictions of all the concentrations are shown in Figures 9.33 and 9.34. Note

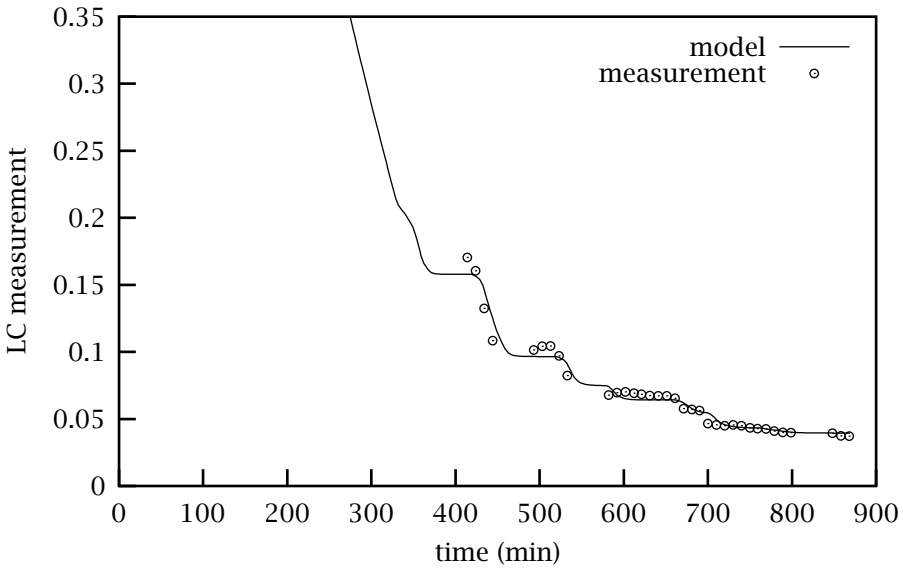

Figure 9.32: Comparison of data to model with optimal parameters.

that the concentration of B provides the most information on the rate constants. Unfortunately this concentration is not measured.

The optimal values of the parameters and their approximate 95% confidence intervals are given by

$$\begin{aligned} n_{A0} &= 2.35 \pm 0.013 \\ k_1 &= 2500 \pm 4800 \\ k_2 &= 1250 \pm 2300 \end{aligned}$$

Notice that the initial amount of A is determined to within 0.6%, but the rate constants have 200% uncertainty. That uncertainty is a direct result of the problem structure. The relative amounts of C and D pin down accurately the amount of starting material. The rate constants determine the speed at which we arrive at these values. After each base addition, we have only rough information about the reaction rates by observing the relative amounts of C and D. If we could measure the B concentration, we could achieve narrow confidence intervals for the rate constants as well. For the purposes of end-point control, however,

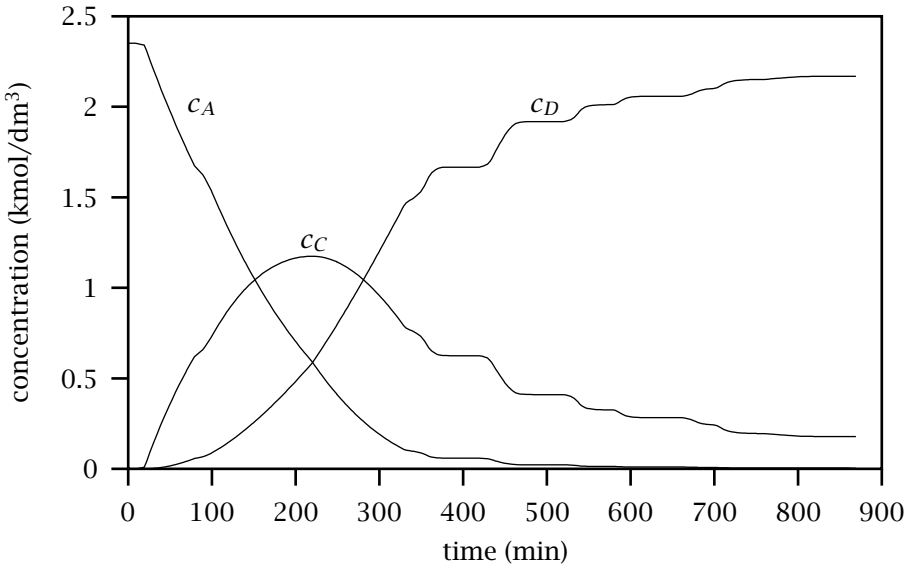

Figure 9.33: Concentrations of species A, C and D versus time.

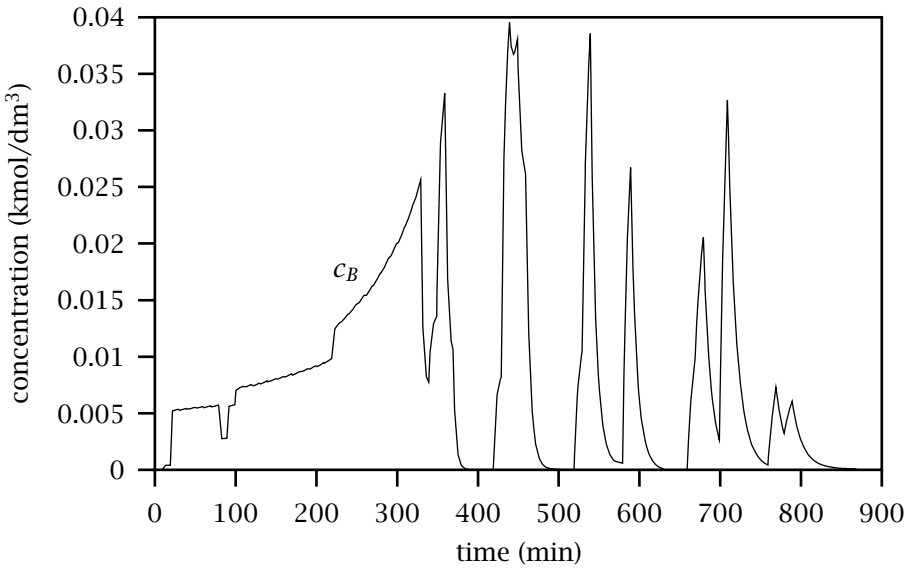

Figure 9.34: Concentration of species B versus time.

the rate constants are irrelevant. We seek to know how much total base to add, not how quickly it will be consumed. In this application, we regard k_1 and k_2 as nuisance parameters. We must estimate them to determine the parameter we care about, but their values are not useful to us.

Model reduction. The large parameter uncertainty tells us that the experimental data do not contain sufficient information to determine the rate constants. That diagnosis is essential because it motivates our next step: model reduction.

Although it may not be immediately apparent, the difficulty we face is caused by the presence of large rate constants. We wish to make the equilibrium assumption as described in Chapter 5 to reduce the model. This analysis is clearest if we first rewrite the material balances in terms of extents of the two reactions as in Chapter 5

$$\begin{aligned}
\frac{dV_R}{dt} &= Q_f \\
\frac{d\varepsilon_1}{dt} &= r_1 = k_1 c_A c_B V_R = k_1 n_A n_B / V_R \\
\frac{d\varepsilon_2}{dt} &= r_2 = k_2 c_C c_B V_R = k_2 n_C n_B / V_R
\end{aligned} \tag{9.72}$$

The initial conditions for this model are

$$\begin{aligned}
V_R(0) &= V_{R0} \\
\varepsilon_1(0) &= 0 \\
\varepsilon_2(0) &= 0
\end{aligned}$$

We can easily translate from the two reaction extents to total moles of species via

$$n_A = n_{A0} - \varepsilon_1 \tag{9.73}$$

$$n_B = n_{B\text{add}} - (\varepsilon_1 + \varepsilon_2) \tag{9.74}$$

$$n_C = \varepsilon_1 - \varepsilon_2 \tag{9.75}$$

$$n_D = \varepsilon_2 \tag{9.76}$$

in which

$$n_{B\text{add}}(t) = \int_0^t Q_f(t') c_{Bf} dt' \tag{9.77}$$

From the two reaction-rate expressions in Equations 9.72, if we assume equilibrium is established with these irreversible reactions, either the

concentration of B is zero or the concentrations of A and C are zero. Due to the reactor-addition policy, we know B is the limiting reagent, and conclude that under large k_1, k_2, the B concentration is zero. The material balance for B in Equation 9.76 then provides one algebraic equation for the two extents

$$\varepsilon_1 + \varepsilon_2 = n_{B\text{add}} \tag{9.78}$$

The additional equation comes from examining the two extents' differential equations and noticing

$$\frac{d\varepsilon_1}{dt} = \left(k\frac{n_A}{n_C}\right)\frac{d\varepsilon_2}{dt} \qquad k = \frac{k_1}{k_2} \tag{9.79}$$

We see that the concentration of B disappears from this slow time-scale model and the ratio $k = k_1/k_2$ appears instead of the individual rate constants. If we wish to cast the reduced model in differential equation form, we can differentiate Equation 9.78 and substitute Equation 9.79 to eliminate ε_1 to obtain

$$\frac{d\varepsilon_2}{dt} = \frac{Q_f c_{Bf}}{1 + kn_A/n_C} \tag{9.80}$$

Substituting Equation 9.76 for the moles of A and C, and using Equation 9.78 again to eliminate ε_1 produces a differential equation for the second extent

$$\frac{d\varepsilon_2}{dt} = Q_f c_{Bf}\left(1 + k\frac{n_{A0} - n_{B\text{add}} + \varepsilon_2}{n_{B\text{add}} - 2\varepsilon_2}\right)^{-1} \tag{9.81}$$

Solving this reduced model for various values of k produces the results shown in Figure 9.35. If we perform parameter estimation with this model and these data we obtain the following results.

$$\boxed{\begin{aligned} n_{A0} &= 2.35 \pm 0.0083 \\ k = k_1/k_2 &= 2.00 \pm 0.29 \end{aligned}} \tag{9.82}$$

We have improved the situation compared to the uncertainty in the full model. But we still have more than 10% uncertainty in the ratio of rate constants k. These results may be adequate, but if we wish to further improve the confidence in $k = k_1/k_2$ and n_{A0}, we proceed by reexamining the experimental design.

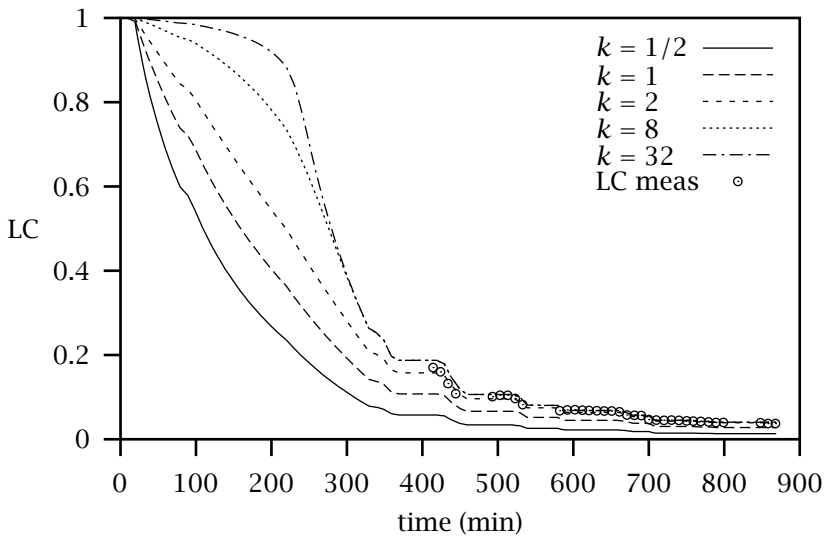

Figure 9.35: Predictions of LC measurement for reduced model.

Experimental design. We have removed the rate constants from the estimation problem by assuming their values are arbitrarily large and making the equilibrium assumption. But we still have a reasonably large uncertainty in the ratio of rate constants. Given what we know at this point, we can easily remedy this final, remaining problem. Consider again the results shown in Figure 9.35. We see that the ratio of the rate constants has its primary effect at early times. By the end of the semi-batch addition, there is little information left. For this reason, we have the relatively large uncertainty in k shown in Equation 9.82. But at early times, when significant amounts of A and C remain in the reactor, the data are much more informative. Although under the original reactor-addition policy, the operator had no reason to turn on the LC until later times, we see that these early measurements are actually the informative ones. We do not have industrial operating data with early LC measurements, but we can simulate the effect with our model. Notice the advantage of modeling. We have the ability to query the model instead of performing expensive experiments to evaluate the impact of a proposed change. Consider Figure 9.36 in which we have simulated early LC measurements by solving the model with

$$k = k_1/k_2 = 2.0, \qquad n_{A0} = 2.35$$

and adding measurement noise. Notice this noise is greater than the

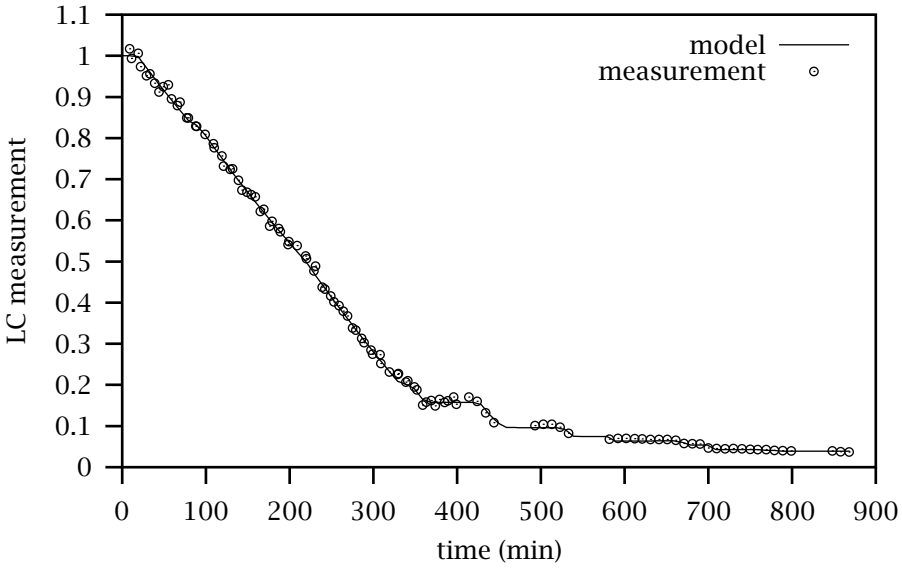

Figure 9.36: Fit of LC measurement versus time for reduced model; uarly time measurements have been added.

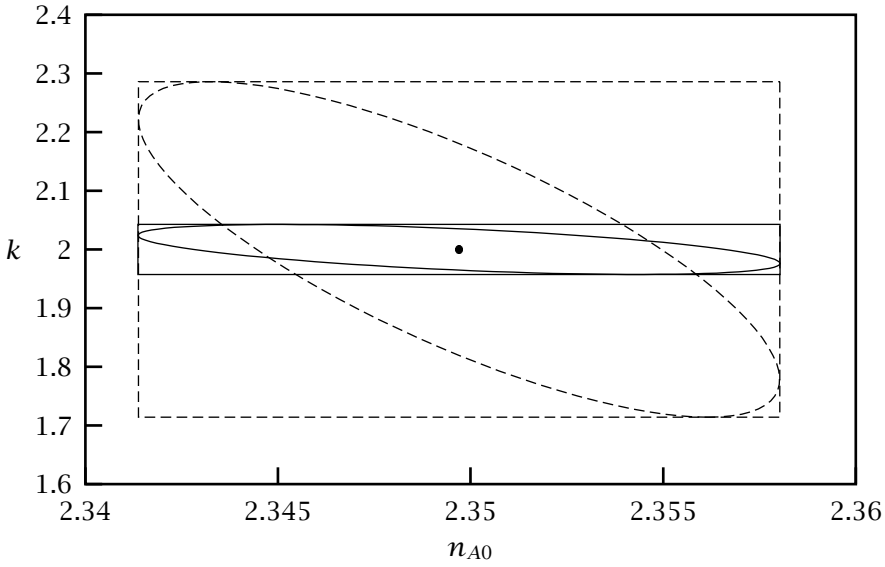

Figure 9.37: Parameter estimates and confidence intervals for reduced model with redesigned experiment.

measurement noise shown in the actual operating data. We next perform parameter estimation on the data shown in Figure 9.36 and obtain the following parameters and approximate confidence intervals

$$\begin{aligned} n_{A0} &= 2.35 \pm 0.0083 \\ k = k_1/k_2 &= 2.00 \pm 0.043 \end{aligned}$$

Figure 9.37 shows the confidence-interval ellipses and boxes when using the original dataset and using the dataset augmented with early-time data. As summarized in Figure 9.37 the primary benefit in adding the early time data is a more precise estimate of k. The uncertainty in k is reduced by more than a factor of six when adding the early time LC measurements. The initial number of moles of A is relatively accurately determined by both datasets.

Improved reactor operation. Given these modeling results, we can shorten the batch time significantly. First we switch the LC on at time $t = 0$. Then as the LC measurements become available, we estimate the initial number of moles n_{A0} and monitor its confidence interval. As soon as the uncertainty in n_{A0} reaches a sufficiently low threshold, we are confident how much B is required and can add the remainder in one shot. Testing this approach with many datasets at Kodak allowed us to conclude that by the time the first large addition was completed, we obtained sufficient confidence on n_{A0} that the rest of the B could be added immediately. Such a procedure reduces the batch time from about 900 minutes with the conventional approach to about half that time with the model-based operation. The new operation essentially doubles the production rate without constructing new reactor facilities, which is significant for this capacity-limited chemical.

9.4 Summary

In this chapter we first summarized some of the analytical methods and experimental reactors used to collect reactor data, focusing the discussion on:

- infrared spectroscopy
- gas chromatography
- mass spectrometry