1 Introduction to Differencing

1.1 A Simple Idea

Consider the nonparametric regression model

$$y = f(x) + \varepsilon \tag{1.1.1}$$

for which little is assumed about the function f except that it is smooth. In its simplest incarnation, the residuals are independently and identically distributed with mean zero and constant variance σ_ε^2, and the x's are generated by a process that ensures they will eventually be dense in the domain. Closeness of the x's combined with smoothness of f provides a basis for estimation of the regression function. By averaging or smoothing observations on y for which the corresponding x's are close to a given point, say x_o, one obtains a reasonable estimate of the regression effect $f(x_o)$.

This premise – that x's that are close will have corresponding values of the regression function that are close – may also be used to remove the regression effect. It is this removal or *differencing* that provides a simple exploratory tool. To illustrate the idea we present four applications:

1. Estimation of the residual variance σ_ε^2,
2. Estimation and inference in the partial linear model $y = z\beta + f(x) + \varepsilon$,
3. A specification test on the regression function f, and
4. A test of equality of nonparametric regression functions.[1]

[1] The first-order differencing estimator of the residual variance in a nonparametric setting appears in Rice (1984). Although unaware of his result at the time, I presented the identical estimator at a conference held at the IC2 Institute at the University of Texas at Austin in May 1984. Differencing subsequently appeared in a series of nonparametric and semiparametric settings, including Powell (1987), Yatchew (1988), Hall, Kay, and Titterington (1990), Yatchew (1997, 1998, 1999, 2000), Lewbel (2000), Fan and Huang (2001), and Horowitz and Spokoiny (2001).

1.2 Estimation of the Residual Variance

Suppose one has data $(y_1, x_1), \ldots, (y_n, x_n)$ on the pure nonparametric regression model (1.1.1), where x is a bounded scalar lying, say, in the unit interval, ε is i.i.d. with $E(\varepsilon \mid x) = 0$, $Var(\varepsilon \mid x) = \sigma_\varepsilon^2$, and all that is known about f is that its first derivative is bounded. Most important, the data have been rearranged so that $x_1 \leq \cdots \leq x_n$. Consider the following estimator of σ_ε^2:

$$s_{diff}^2 = \frac{1}{2n} \sum_{i=2}^{n} (y_i - y_{i-1})^2. \tag{1.2.1}$$

The estimator is consistent because, as the x's become close, differencing tends to remove the nonparametric effect $y_i - y_{i-1} = f(x_i) - f(x_{i-1}) + \varepsilon_i - \varepsilon_{i-1} \cong \varepsilon_i - \varepsilon_{i-1}$, so that[2]

$$s_{diff}^2 \cong \frac{1}{2n} \sum_{i=2}^{n} (\varepsilon_i - \varepsilon_{i-1})^2 \cong \frac{1}{n} \sum_{i=1}^{n} \varepsilon_i^2 - \frac{1}{n} \sum_{i=2}^{n} \varepsilon_i \varepsilon_{i-1}. \tag{1.2.2}$$

An obvious advantage of s_{diff}^2 is that no initial estimate of the regression function f needs to be calculated. Indeed, no consistent estimate of f is implicit in (1.2.1). Nevertheless, the terms in s_{diff}^2 that involve f converge to zero sufficiently quickly so that the asymptotic distribution of the estimator can be derived directly from the approximation in (1.2.2). In particular,

$$n^{1/2}\left(s_{diff}^2 - \sigma_\varepsilon^2\right) \xrightarrow{D} N(0, E(\varepsilon^4)). \tag{1.2.3}$$

Moreover, derivation of this result is facilitated by the assumption that the ε_i are independent so that reordering of the data does not affect the distribution of the right-hand side in (1.2.2).

1.3 The Partial Linear Model

Consider now the partial linear model $y = z\beta + f(x) + \varepsilon$, where for simplicity all variables are assumed to be scalars. We assume that $E(\varepsilon \mid z, x) = 0$ and that $Var(\varepsilon \mid z, x) = \sigma_\varepsilon^2$.[3] As before, the x's have bounded support, say the unit interval, and have been rearranged so that $x_1 \leq \cdots \leq x_n$. Suppose that the conditional mean of z is a smooth function of x, say $E(z \mid x) = g(x)$ where g' is

[2] To see why this approximation works, suppose that the x_i are equally spaced on the unit interval and that $f' \leq L$. By the mean value theorem, for some $x_i^* \in [x_{i-1}, x_i]$ we have $f(x_i) - f(x_{i-1}) = f'(x_i^*)(x_i - x_{i-1}) \leq L/n$. Thus, $y_i - y_{i-1} = \varepsilon_i - \varepsilon_{i-1} + O(1/n)$. For detailed development of the argument, see Exercise 1. If the x_i have a density function bounded away from zero on the support, then $x_i - x_{i-1} \cong O_P(1/n)$ and $y_i - y_{i-1} \cong \varepsilon_i - \varepsilon_{i-1} + O_P(1/n)$. See Appendix B, Lemma B.2, for a related result.

[3] For extensions to the heteroskedastic and autocorrelated cases, see Sections 3.6 and 4.5.

bounded and $Var(z \mid x) = \sigma_u^2$. Then we may rewrite $z = g(x) + u$. Differencing yields

$$\begin{aligned} y_i - y_{i-1} &= (z_i - z_{i-1})\beta + (f(x_i) - f(x_{i-1})) + \varepsilon_i - \varepsilon_{i-1} \\ &= (g(x_i) - g(x_{i-1}))\beta + (u_i - u_{i-1})\beta \\ &\quad + (f(x_i) - f(x_{i-1})) + \varepsilon_i - \varepsilon_{i-1} \\ &\cong (u_i - u_{i-1})\beta + \varepsilon_i - \varepsilon_{i-1}. \end{aligned} \tag{1.3.1}$$

Thus, the direct effect $f(x)$ of the nonparametric variable x and the indirect effect $g(x)$ that occurs through z are removed. Suppose we apply the OLS estimator of β to the differenced data, that is,

$$\hat{\beta}_{diff} = \frac{\sum(y_i - y_{i-1})(z_i - z_{i-1})}{\sum(z_i - z_{i-1})^2}. \tag{1.3.2}$$

Then, substituting the approximations $z_i - z_{i-1} \cong u_i - u_{i-1}$ and $y_i - y_{i-1} \cong (u_i - u_{i-1})\beta + \varepsilon_i - \varepsilon_{i-1}$ into (1.3.2) and rearranging, we have

$$n^{1/2}(\hat{\beta}_{diff} - \beta) \cong \frac{n^{1/2}\frac{1}{n}\sum(\varepsilon_i - \varepsilon_{i-1})(u_i - u_{i-1})}{\frac{1}{n}\sum(u_i - u_{i-1})^2}. \tag{1.3.3}$$

The denominator converges to $2\,\sigma_u^2$, and the numerator has mean zero and variance $6\,\sigma_\varepsilon^2\sigma_u^2$. Thus, the ratio has mean zero and variance $6\,\sigma_\varepsilon^2\sigma_u^2/(2\sigma_u^2)^2 = 1.5\,\sigma_\varepsilon^2/\sigma_u^2$. Furthermore, the ratio may be shown to be approximately normal (using a finitely dependent central limit theorem). Thus, we have

$$n^{1/2}(\hat{\beta}_{diff} - \beta) \xrightarrow{D} N\left(0, \frac{1.5\,\sigma_\varepsilon^2}{\sigma_u^2}\right). \tag{1.3.4}$$

For the most efficient estimator, the corresponding variance in (1.3.4) would be $\sigma_\varepsilon^2/\sigma_u^2$ so the proposed estimator based on first differences has relative efficiency $2/3 = 1/1.5$. In Chapters 3 and 4 we will produce efficient estimators.

Now, in order to use (1.3.4) to perform inference, we will need consistent estimators of σ_ε^2 and σ_u^2. These may be obtained using

$$\begin{aligned} s_\varepsilon^2 &= \frac{1}{2n}\sum_{i=2}^{n}((y_i - y_{i-1}) - (z_i - z_{i-1})\hat{\beta}_{diff})^2 \\ &\cong \frac{1}{2n}\sum_{i=2}^{n}(\varepsilon_i - \varepsilon_{i-1})^2 \xrightarrow{P} \sigma_\varepsilon^2 \end{aligned} \tag{1.3.5}$$

and

$$s_u^2 = \frac{1}{2n}\sum_{i=2}^{n}(z_i - z_{i-1})^2 \cong \frac{1}{2n}\sum_{i=2}^{n}(u_i - u_{i-1})^2 \xrightarrow{P} \sigma_u^2. \tag{1.3.6}$$

The preceding procedure generalizes straightforwardly to models with multiple parametric explanatory variables.

1.4 Specification Test

Suppose, for example, one wants to test the null hypothesis that f is a linear function. Let s^2_{res} be the usual estimate of the residual variance obtained from a linear regression of y on x. If the linear model is correct, then s^2_{res} will be approximately equal to the average of the true squared residuals:

$$s^2_{res} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{\gamma}_1 - \hat{\gamma}_2 x_i)^2 \cong \frac{1}{n}\sum_{i=1}^{n}\varepsilon_i^2. \tag{1.4.1}$$

If the linear specification is incorrect, then s^2_{res} will overestimate the residual variance while s^2_{diff} in (1.2.1) will remain a consistent estimator, thus forming the basis of a test. Consider the test statistic

$$V = \frac{n^{1/2}\left(s^2_{res} - s^2_{diff}\right)}{s^2_{diff}}. \tag{1.4.2}$$

Equations (1.2.2) and (1.4.1) imply that the numerator of V is approximately equal to

$$n^{1/2}\frac{1}{n}\sum \varepsilon_i\,\varepsilon_{i-1} \xrightarrow{D} N\left(0, \sigma_\varepsilon^4\right). \tag{1.4.3}$$

Since s^2_{diff}, the denominator of V, is a consistent estimator of σ^2_ε, V is asymptotically $N(0,1)$ under H_0. (Note that this is a one-sided test, and one rejects for large values of the statistic.)

As we will see later, this test procedure may be used to test a variety of null hypotheses such as general parametric and semiparametric specifications, monotonicity, concavity, additive separability, and other constraints. One simply inserts the restricted estimator of the variance in (1.4.2). We refer to test statistics that compare restricted and unrestricted estimates of the residual variance as "goodness-of-fit" tests.

1.5 Test of Equality of Regression Functions

Suppose we are given data $(y_{A1}, x_{A1}), \ldots, (y_{An}, x_{An})$ and $(y_{B1}, x_{B1}), \ldots, (y_{Bn}, x_{Bn})$ from two possibly different regression models A and B. Assume x is a scalar and that each data set has been reordered so that the x's are in increasing order. The basic models are

$$\begin{aligned} y_{Ai} &= f_A(x_{Ai}) + \varepsilon_{Ai} \\ y_{Bi} &= f_B(x_{Bi}) + \varepsilon_{Bi} \end{aligned} \tag{1.5.1}$$

where given the x's, the ε's have mean 0, variance σ_ε^2, and are independent within and between populations; f_A and f_B have first derivatives bounded. Using (1.2.1), define consistent "within" differencing estimators of the variance

$$\begin{aligned} s_A^2 &= \frac{1}{2n}\sum_i^n (y_{Ai} - y_{Ai-1})^2 \\ s_B^2 &= \frac{1}{2n}\sum_i^n (y_{Bi} - y_{Bi-1})^2. \end{aligned} \tag{1.5.2}$$

As we will do frequently, we have dropped the subscript "*diff*". Now pool *all* the data and reorder so that the pooled x's are in increasing order: $(y_1^*, x_1^*), \ldots\ldots, (y_{2n}^*, x_{2n}^*)$. (Note that the pooled data have only one subscript.) Applying the differencing estimator once again, we have

$$s_p^2 = \frac{1}{4n}\sum_j^{2n} \left(y_j^* - y_{j-1}^*\right)^2. \tag{1.5.3}$$

The basic idea behind the test procedure is to compare the pooled estimator with the average of the within estimators. If $f_A = f_B$, then the within and pooled estimators are consistent and should yield similar estimates. If $f_A \neq f_B$, then the within estimators remain consistent, whereas the pooled estimator overestimates the residual variance, as may be seen in Figure 1.1.

To formalize this idea, define the test statistic

$$\Upsilon \equiv (2n)^{1/2}\left(s_p^2 - 1/2\left(s_A^2 + s_B^2\right)\right). \tag{1.5.4}$$

If $f_A = f_B$, then differencing removes the regression effect sufficiently quickly in both the within and the pooled estimators so that

$$\begin{aligned} \Upsilon &\equiv (2n)^{1/2}\left(s_p^2 - 1/2\left(s_A^2 + s_B^2\right)\right) \\ &\cong \frac{(2n)^{1/2}}{4n}\left(\sum_j^{2n}\left(\varepsilon_j^* - \varepsilon_{j-1}^*\right)^2 - \sum_i^n (\varepsilon_{Ai} - \varepsilon_{Ai-1})^2 - \sum_i^n (\varepsilon_{Bi} - \varepsilon_{Bi-1})^2\right) \\ &\cong \frac{(2n)^{1/2}}{2n}\left(\sum_j^{2n}\varepsilon_j^{*2} - \varepsilon_j^*\varepsilon_{j-1}^* - \sum_i^n \varepsilon_{Ai}^2 - \varepsilon_{Ai}\varepsilon_{Ai-1} - \sum_i^n \varepsilon_{Bi}^2 - \varepsilon_{Bi}\varepsilon_{Bi-1}\right) \\ &\cong \frac{1}{(2n)^{1/2}}\left(\sum_i^n \varepsilon_{Ai}\varepsilon_{Ai-1} + \sum_i^n \varepsilon_{Bi}\varepsilon_{Bi-1}\right) - \frac{1}{(2n)^{1/2}}\left(\sum_j^{2n}\varepsilon_j^*\varepsilon_{j-1}^*\right). \end{aligned} \tag{1.5.5}$$

Consider the two terms in the last line. In large samples, each is approximately $N(0, \sigma_\varepsilon^4)$. If observations that are consecutive in the individual data

Within estimators of residual variance

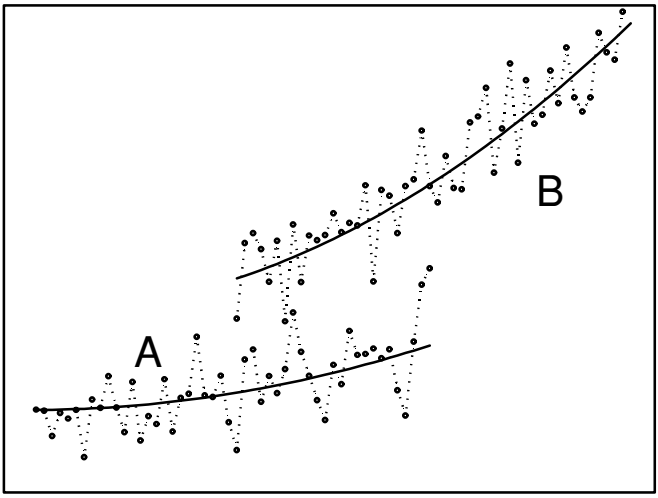

Pooled estimator of residual variance

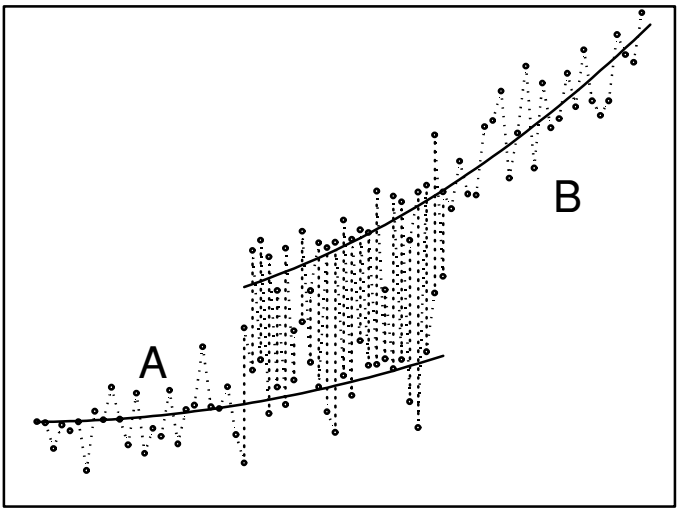

Figure 1.1. Testing equality of regression functions.

sets tend to be consecutive after pooling and reordering, then the *covariance* between the two terms will be large. In particular, the covariance is approximately $\sigma_\varepsilon^4(1-\pi)$, where π equals the probability that consecutive observations in the pooled reordered data set come from *different* populations.

It follows that under $H_o : f_A = f_B$,

$$\Upsilon \xrightarrow{D} N\left(0, 2\pi\sigma_\varepsilon^4\right). \tag{1.5.6}$$

For example, if reordering the pooled data is equivalent to stacking data sets A and B – because the two sets of x's, x_A and x_B, do not intersect – then $\pi \cong 0$ and indeed the statistic Υ becomes degenerate. This is not surprising, since observing nonparametric functions over different domains cannot provide a basis for testing whether they are the same. If the pooled data involve a simple interleaving of data sets A and B, then $\pi \cong 1$ and $\Upsilon \to N(0, 2\sigma_\varepsilon^4)$. If x_A and x_B are independent of each other but have the same distribution, then for the pooled reordered data the probability that consecutive observations come from different populations is $1/2$ and $\Upsilon \to N(0, \sigma_\varepsilon^4)$.[4] To implement the test, one may obtain a consistent estimate $\hat{\pi}$ by taking the proportion of observations in the pooled reordered data that are preceded by an observation from a different population.

1.6 Empirical Application: Scale Economies in Electricity Distribution[5]

To illustrate these ideas, consider a simple variant of the Cobb–Douglas model for the costs of distributing electricity

$$\begin{aligned} tc = {} & f(cust) + \beta_1 wage + \beta_2 pcap \\ & + \beta_3 PUC + \beta_4 kwh + \beta_5 life + \beta_6 lf + \beta_7 kmwire + \varepsilon \end{aligned} \tag{1.6.1}$$

where *tc* is the log of total cost per customer, *cust* is the log of the number of customers, *wage* is the log wage rate, *pcap* is the log price of capital, *PUC* is a dummy variable for public utility commissions that deliver additional services and therefore may benefit from economies of scope, *life* is the log of the remaining life of distribution assets, *lf* is the log of the load factor (this measures capacity utilization relative to peak usage), and *kmwire* is the log of kilometers of distribution wire per customer. The data consist of 81 municipal distributors in Ontario, Canada, during 1993. (For more details, see Yatchew, 2000.)

[4] For example, distribute n men and n women randomly along a stretch of beach facing the sunset. Then, for any individual, the probability that the person to the left is of the opposite sex is $1/2$. More generally, if x_A and x_B are independent of each other and have different distributions, then π depends on the relative density of observations from each of the two populations.

[5] Variable definitions for empirical examples are contained in Appendix E.

Because the data have been reordered so that the nonparametric variable *cust* is in increasing order, first differencing (1.6.1) tends to remove the nonparametric effect f. We also divide by $\sqrt{2}$ so that the residuals in the differenced Equation (1.6.2) have the same variance as those in (1.6.1). Thus, we have

$$\begin{aligned}
&[tc_i - tc_{i-1}]/\sqrt{2} \\
&\quad \cong \beta_1[wage_i - wage_{i-1}]/\sqrt{2} + \beta_2[pcap_i - pcap_{i-1}]/\sqrt{2} \\
&\qquad + \beta_3[PUC_i - PUC_{i-1}]/\sqrt{2} + \beta_4[kwh_i - kwh_{i-1}]/\sqrt{2} \\
&\qquad + \beta_5[life_i - life_{i-1}]/\sqrt{2} + \beta_6[lf_i - lf_{i-1}]/\sqrt{2} \\
&\qquad + \beta_7[kmwire_i - kmwire_{i-1}]/\sqrt{2} + [\varepsilon_i - \varepsilon_{i-1}]/\sqrt{2}. \qquad (1.6.2)
\end{aligned}$$

Figure 1.2 summarizes our estimates of the parametric effects β using the differenced equation. It also contains estimates of a pure parametric specification in which the scale effect f is modeled with a quadratic. Applying the specification test (1.4.2), where s^2_{diff} is replaced with (1.3.5), yields a value of 1.50, indicating that the quadratic model may be adequate.

Thus far our results suggest that by differencing we can perform inference on β as if there were no nonparametric component f in the model to begin with. But, having estimated β, we can then proceed to apply a variety of nonparametric techniques to analyze f as if β were known. Such a modular approach simplifies implementation because it permits the use of existing software designed for pure nonparametric models.

More precisely, suppose we assemble the ordered pairs $(y_i - z_i\hat{\beta}_{diff}, x_i)$; then, we have

$$y_i - z_i\hat{\beta}_{diff} = z_i(\beta - \hat{\beta}_{diff}) + f(x_i) + \varepsilon_i \cong f(x_i) + \varepsilon_i. \qquad (1.6.3)$$

If we apply conventional smoothing methods to these ordered pairs such as kernel estimation (see Section 3.2), then consistency, optimal rate of convergence results, and the construction of confidence intervals for f remain valid because $\hat{\beta}_{diff}$ converges sufficiently quickly to β that the approximation in the last part of (1.6.3) leaves asymptotic arguments unaffected. (This is indeed why we could apply the specification test after removing the *estimated* parametric effect.) Thus, in Figure 1.2 we have also plotted a nonparametric (kernel) estimate of f that can be compared with the quadratic estimate. In subsequent sections, we will elaborate this example further and provide additional ones.

1.7 Why Differencing?

An important advantage of differencing procedures is their simplicity. Consider once again the partial linear model $y = z\beta + f(x) + \varepsilon$. Conventional

Variable	Quadratic model		Partial linear model[a]	
	Coef	SE	Coef	SE
cust	−0.833	0.175	–	–
$cust^2$	0.040	0.009	–	–
wage	0.833	0.325	0.448	0.367
pcap	0.562	0.075	0.459	0.076
PUC	−0.071	0.039	−0.086	0.043
kwh	−0.017	0.089	−0.011	0.087
life	−0.603	0.119	−0.506	0.131
lf	1.244	0.434	1.252	0.457
kmwire	0.445	0.086	0.352	0.094
s^2_ε	.021		.018	
R^2	.618		.675	

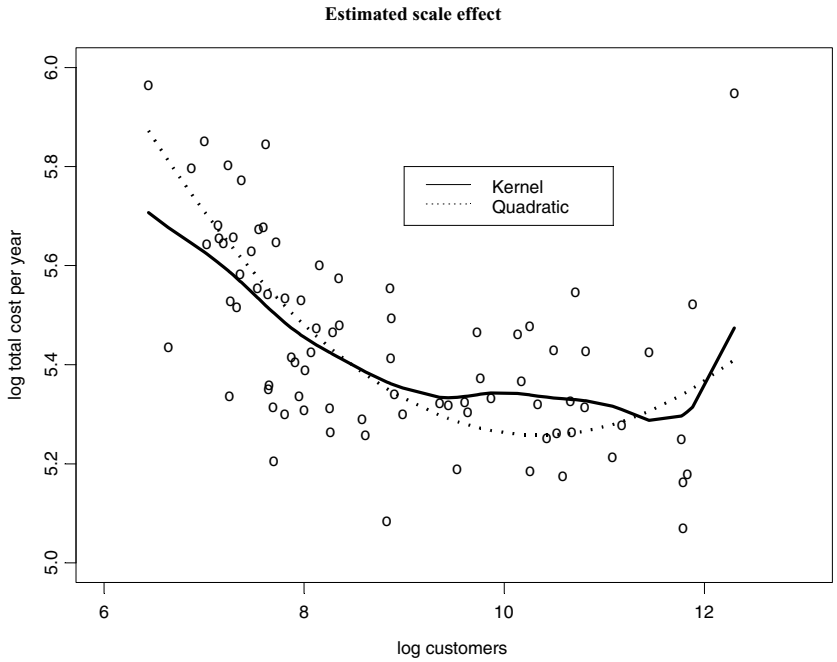

[a] Test of quadratic versus nonparametric specification of scale effect: $V = n^{1/2}(s^2_{res} - s^2_{diff})/s^2_{diff} = 81^{1/2}(.021 - .018)/.018 = 1.5$, where V is $N(0,1)$, Section 1.4.

Figure 1.2. Partial linear model – Log-linear cost function: Scale economies in electricity distribution.

estimators, such as the one proposed by Robinson (1988) (see Section 3.6), require one to estimate $E(y \mid x)$ and $E(z \mid x)$ using nonparametric regressions. The estimated residuals from each of these regressions (hence the term "double residual method") are then used to estimate the *parametric* regression

$$y - E(y \mid x) = (z - E(z \mid x))\beta + \varepsilon. \tag{1.7.1}$$

If z is a vector, then a separate nonparametric regression is run for each component of z, where the independent variable is the nonparametric variable x. In contrast, differencing eliminates these first-stage regressions so that estimation of β can be performed – regardless of its dimension – even if nonparametric regression procedures are not available within the software being used. Similarly, tests of parametric specifications against nonparametric alternatives and tests of equality of regression functions across two or more (sub-) samples can be carried out without performing a nonparametric regression.

As should be evident from the empirical example of the last section, differencing may easily be combined with other procedures. In that example, we used differencing to estimate the parametric component of a partial linear model. We then removed the estimated parametric effect and applied conventional nonparametric procedures to analyze the nonparametric component. Such modular analysis does require theoretical justification, which we will provide in Section 4.12.

As we have seen, the partial linear model permits a simple semiparametric generalization of the Cobb–Douglas model. Translog and other linear-in-parameters models may be generalized similarly. If we allow the parametric portion of the model to be nonlinear – so that we have a partial parametric model – then we may also obtain simple semiparametric generalizations of models such as the constant elasticity of substitution (CES) cost function. These, too, may be estimated straightforwardly using differencing (see Section 4.7). The key requirement is that the parametric and nonparametric portions of the model be additively separable.

Other procedures commonly used by the econometrician may be imported into the differencing setting with relative ease. If some of the parametric variables are potentially correlated with the residuals, instrumental variable techniques can be applied, with suitable modification, as can the Hausman endogeneity test (see Section 4.8). If the residuals are potentially not homoskedastic, then well-known techniques such as White's heteroskedasticity-consistent standard errors can be adapted (see Section 4.5). The reader will no doubt find other procedures that can be readily transplanted.

Earlier we have pointed out that the first-order differencing estimator of β in the partial linear model is inefficient when compared with the most efficient estimator (see Section 1.3). The same is true for the first-order differencing estimator of the residual variance (see Section 1.2). This problem can be corrected using higher-order differencing, as demonstrated in Chapter 4.