
Testing Parametric Regression Specifications with Nonparametric Regression

9.1 Testing Functional Forms

An important, if under-appreciated, use of nonparametric regression is checking whether parametric regressions are well-specified. The typical parametric regression model is something like

$$Y = f(X; \theta) + \epsilon \quad (9.1)$$

where f is some function which is completely specified except for the finite vector of parameters θ , and ϵ , as usual, is uncorrelated noise. Often, of course, people use a function f that is linear in the variables in X , or perhaps includes some interactions between them.

How can we tell if the specification is right? If, for example, it's a linear model, how can we check whether there might not be some nonlinearity? A common approach is to modify the specification to allow for *specific* departures from the baseline model — say, adding a quadratic term — and seeing whether the coefficients that go with those terms are significantly non-zero, or whether the improvement in fit is significant.¹ For example, one might compare the model

$$Y = \theta_1 x_1 + \theta_2 x_2 + \epsilon \quad (9.2)$$

to the model

$$Y = \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_1^2 + \epsilon \quad (9.3)$$

by checking whether the estimated θ_3 is significantly different from 0, or whether the residuals from the second model are significantly smaller than the residuals from the first.

This can work, *if* you have chosen the right nonlinearity to test. It has the power to detect certain mis-specifications, if they exist, but not others. (What if the departure from linearity is not quadratic but cubic?) If you have good reasons to think that when the model is wrong, it can only be wrong in certain ways, fine; if not, though, why only check for those errors?

Nonparametric regression effectively lets you check for *all* kinds of systematic errors, rather than singling out a particular one. There are three basic approaches, which I give in order of increasing sophistication.

¹ In my experience, this approach is second in popularity only to ignoring the issue.

- If the parametric model is right, it should predict as well as, or even better than, the non-parametric one, and we can check whether $MSE_p(\hat{\theta}) - MSE_{np}(\hat{\mu})$ is sufficiently small.
- If the parametric model is right, the non-parametric estimated regression curve should be very close to the parametric one. So we can check whether $f(x; \hat{\theta}) - \hat{\mu}(x)$ is approximately zero everywhere.
- If the parametric model is right, then its *residuals* should be patternless and independent of input features, because

$$\mathbb{E}[Y - f(x; \theta)|X] = \mathbb{E}[f(x; \theta) + \epsilon - f(x; \theta)|X] = \mathbb{E}[\epsilon|X] = 0 \quad (9.4)$$

So we can apply non-parametric smoothing to the parametric residuals, $y - f(x; \hat{\theta})$, and see if their expectation is approximately zero everywhere.

We'll stick with the first procedure, because it's simpler for us to implement computationally. However, it turns out to be easier to develop theory for the other two, and especially for the third — see [Li and Racine \(2007 ch. 12\)](#), or [Hart \(1997\)](#).

Here is the basic procedure.

1. Get data $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$.
2. Fit the parametric model, getting an estimate $\hat{\theta}$, and in-sample mean-squared error $MSE_p(\hat{\theta})$.
3. Fit your favorite nonparametric regression (using cross-validation to pick control settings as necessary), getting curve $\hat{\mu}$ and in-sample mean-squared error $MSE_{np}(\hat{\mu})$.
4. Calculate $\hat{d} = MSE_p(\hat{\theta}) - MSE_{np}(\hat{\mu})$.
5. Simulate from the parametric model $\hat{\theta}$ to get faked data $(x_1^*, y_1^*), \dots, (x_n^*, y_n^*)$.
 1. Fit the parametric model to the simulated data, getting estimate $\hat{\theta}^*$ and $MSE_p(\hat{\theta}^*)$.
 2. Fit the nonparametric model to the simulated data, getting estimate $\hat{\mu}^*$ and $MSE_{np}(\hat{\mu}^*)$.
 3. Calculate $D^* = MSE_p(\hat{\theta}^*) - MSE_{np}(\hat{\mu}^*)$.
6. Repeat step [5](#) b times to get an estimate of the distribution of D under the null hypothesis.
7. The approximate p -value is $\frac{1 + \#\{D^* > \hat{d}\}}{1 + b}$.

Let's step through the logic. In general, the error of the non-parametric model will be converging to the smallest level compatible with the intrinsic noise of the process. What about the parametric model?

Suppose on the one hand that the parametric model *is* correctly specified. Then its error will also be converging to the minimum — by assumption, it's got the functional form right so bias will go to zero, and as $\hat{\theta} \rightarrow \theta_0$, the variance will also

go to zero. In fact, with enough data the correctly-specified parametric model will actually *generalize* better than the non-parametric model.²

Suppose on the other hand that the parametric model is mis-specified. Then its predictions are systematically wrong, even with unlimited amounts of data — there's some bias which never goes away, no matter how big the sample. Since the non-parametric smoother does eventually come arbitrarily close to the true regression function, the smoother will end up predicting better than the parametric model.

Smaller errors for the smoother, then, suggest that the parametric model is wrong. But since the smoother has higher capacity, it could easily get smaller errors on a particular sample by chance and/or over-fitting, so only *big* differences in error count as evidence. Simulating from the parametric model gives us surrogate data which looks just like reality ought to, *if* the model is true. We then see how much better we could expect the non-parametric smoother to fit *under the parametric model*. If the non-parametric smoother fits the actual data much better than this, we can reject the parametric model with high confidence: it's really unlikely that we'd see that big an improvement from using the nonparametric model just by luck.³

As usual, we simulate from the parametric model simply because we have no hope of working out the distribution of the differences in MSEs from first principles. This is an example of our general strategy of bootstrapping.

² Remember that the smoother must, so to speak, use up some of the information in the data to figure out the shape of the regression function. The parametric model, on the other hand, takes that basic shape as given, and uses all the data's information to tune its parameters.

³ As usual with p -values, this is not symmetric. A high p -value might mean that the true regression function is very close to $\mu(x; \theta)$, or it might mean that we don't have enough data to draw conclusions (or that we were unlucky).

9.1.1 Examples of Testing a Parametric Model

Let's see this in action. First, let's detect a reasonably subtle nonlinearity. Take the non-linear function $g(x) = \log(1 + x)$, and say that $Y = g(x) + \epsilon$, with ϵ being IID Gaussian noise with mean 0 and standard deviation 0.15. (This is one of the examples from [4.2]) Figure [9.1] shows the regression function and the data. The nonlinearity is clear with the curve to “guide the eye”, but fairly subtle.

A simple linear regression looks pretty good:

```
glinfit = lm(y ~ x, data = gframe)
print(summary(glinfit), signif.stars = FALSE, digits = 2)
##
## Call:
## lm(formula = y ~ x, data = gframe)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -0.468 -0.109 -0.002  0.114  0.485
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    0.218      0.021     11 <2e-16
## x              0.420      0.012     36 <2e-16
##
## Residual standard error: 0.18 on 298 degrees of freedom
## Multiple R-squared:  0.81, Adjusted R-squared:  0.81
## F-statistic: 1.3e+03 on 1 and 298 DF,  p-value: <2e-16
```

R^2 is ridiculously high — the regression line preserves 81 percent of the variance in the data. The p -value reported by R is also very, very low, but remember all this really means is “you’d have to be crazy to think a flat line fit better than straight line with a slope” (Figure [9.2]).

The in-sample MSE of the linear fit is⁴

```
signif(mean(residuals(glinfit)^2), 3)
## [1] 0.0311
```

The nonparametric regression has a somewhat smaller MSE⁵


```
library(np)
gnpr <- npreg(y ~ x, data = gframe)
signif(gnpr$MSE, 3)
## [1] 0.024
```

So \hat{d} is

```
signif((d.hat = mean(glinfit$residual^2) - gnpr$MSE), 3)
## [1] 0.00715
```

⁴ If we ask R for the MSE, by squaring `summary(glinfit)$sigma`, we get 0.031346. This differs from the mean of the squared residuals by a factor of factor of $n/(n-2) = 300/298 = 1.0067$, because R is trying to estimate the out-of-sample error by scaling up the in-sample error, the same way the estimated population variance scales up the sample variance. We want to compare in-sample fits.

⁵ `npreg` does not apply the kind of correction mentioned in the previous footnote.



```
x <- runif(300, 0, 3)
yg <- log(x + 1) + rnorm(length(x), 0, 0.15)
gframe <- data.frame(x = x, y = yg)
plot(x, yg, xlab = "x", ylab = "y", pch = 16, cex = 0.5)
curve(log(1 + x), col = "grey", add = TRUE, lwd = 4)
```

Figure 9.1 True regression curve (grey) and data points (circles). The curve $g(x) = \log(1 + x)$.

Now we need to simulate from the fitted parametric model, using its estimated coefficients and noise level. We have seen several times now how to do this. The function `sim.lm` in Example [21](#) does this, along the same lines as the examples in Chapter [6](#) it assumes homoskedastic Gaussian noise. Again, as before, we need




Figure 9.2 As in previous figure, but adding the least-squares regression line (black). Line widths exaggerated for clarity.

a function which will calculate the difference in MSEs between a linear model and a kernel smoother fit to the same data set — which will do automatically what we did by hand above. This is `calc.D` in Example 22. Note that the kernel bandwidth has to be re-tuned to each new data set.

If we call `calc.D` on the output of `sim.lm`, we get one value of the test statistic under the null distribution:

```
calc.D(sim.lm(glinfit, x))
## [1] 0.002131516
```

Now we just repeat this a lot to get a good approximation to the sampling distribution of D under the null hypothesis:

```

sim.lm <- function(linfit, test.x) {
  n <- length(test.x)
  sim.frame <- data.frame(x = test.x)
  sigma <- summary(linfit)$sigma * (n - 2)/n
  y.sim <- predict(linfit, newdata = sim.frame)
  y.sim <- y.sim + rnorm(n, 0, sigma)
  sim.frame <- data.frame(sim.frame, y = y.sim)
  return(sim.frame)
}

```

CODE EXAMPLE 21: *Simulate a new data set from a linear model, assuming homoskedastic Gaussian noise. It also assumes that there is one input variable, x , and that the response variable is called y . Could you modify it to work with multiple regression?*

```

calc.D <- function(data) {
  MSE.p <- mean((lm(y ~ x, data = data)$residuals)^2)
  MSE.np.bw <- npregbw(y ~ x, data = data)
  MSE.np <- npreg(MSE.np.bw)$MSE
  return(MSE.p - MSE.np)
}

```

CODE EXAMPLE 22: *Calculate the difference-in-MSEs test statistic.*

```
null.samples.D <- replicate(200, calc.D(sim.lm(glinfit, x)))
```

This takes some time, because each replication involves not just generating a new simulation sample, but also cross-validation to pick a bandwidth. This adds up to about a second per replicate on my laptop, and so a couple of minutes for 200 replicates.

(While the computer is thinking, look at the command a little more closely. It leaves the x values alone, and only uses simulation to generate new y values. This is appropriate here because our model doesn't really *say* where the x values came from; it's just about the conditional distribution of Y *given* X . If the model we were testing specified a distribution for x , we should generate x each time we invoke `calc.D`. If the specification is vague, like “ x is IID” but with no particular distribution, then resample X .)

When it's done, we can plot the distribution and see that the observed value \hat{d} is pretty far out along the right tail (Figure 9.3). This tells us that it's very unlikely that `npreg` would improve so much on the linear model if the latter were true. In fact, exactly 0 of the simulated values of the test statistic were that big:

```


sum(null.samples.D > d.hat)
## [1] 0

```

Thus our estimated p -value is ≤ 0.00498 . We can reject the linear model pretty confidently⁶

As a second example, let's suppose that the linear model *is* right — then the

⁶ If we wanted a more precise estimate of the p -value, we'd need to use more bootstrap samples.



```
hist(null.samples.D, n = 31, xlim = c(min(null.samples.D), 1.1 * d.hat), probability = TRUE)
abline(v = d.hat)
```


Figure 9.3 Histogram of the distribution of $D = MSE_p - MSE_{np}$ for data simulated from the parametric model. The vertical line marks the observed value. Notice that the mode is positive and the distribution is right-skewed; this is typical.

test should give us a high p -value. So let us stipulate that in reality

$$Y = 0.2 + 0.5x + \eta \quad (9.5)$$

with $\eta \sim \mathcal{N}(0, 0.15^2)$. Figure 9.4 shows data from this, of the same size as before.

Repeating the same exercise as before, we get that $\hat{d} = 0.0013$, together with



```
y2 <- 0.2 + 0.5 * x + rnorm(length(x), 0, 0.15)
y2.frame <- data.frame(x = x, y = y2)
plot(x, y2, xlab = "x", ylab = "y")
abline(0.2, 0.5, col = "grey", lwd = 2)
```

Figure 9.4 Data from the linear model (true regression line in grey).

a slightly different null distribution (Figure [9.5](#)). Now the p -value is 0.14, which it would be quite rash to reject.




Figure 9.5 As in Figure 9.3 but using the data and fits from Figure 9.4

9.1.2 Remarks

Other Nonparametric Regressions

There is nothing especially magical about using kernel regression here. Any consistent nonparametric estimator (say, your favorite spline) would work. They may differ somewhat in their answers on particular cases.

Curse of Dimensionality

For multivariate regressions, testing against a fully nonparametric alternative can be very time-consuming, as well as running up against curse-of-dimensionality issues⁷. A compromise is to test the parametric regression against an additive model. Essentially nothing has to change.

Testing $\mathbb{E}[\hat{\epsilon}|X] = 0$

I mentioned at the beginning of the chapter that one way to test whether the parametric model is correctly specified is to test whether the residuals have expectation zero everywhere. Setting $r(x; m) \equiv \mathbb{E}[Y - m(X)|X = x]$, we know from Chapter ?? that $r(x; \mu) = 0$ everywhere, and that, for any other function m , $r(x; m) \neq 0$ for at least some values of x . Thus, if we take the residuals $\hat{\epsilon}$ from our parametric model and we smooth them, we get an estimated function $\hat{r}(x)$ that should be converging to 0 everywhere if the parametric model is well-specified. A natural test statistic is therefore some measure of the “size” of \hat{r} , such as⁸ $\int \hat{r}^2(x)dx$, or $\int \hat{r}^2(x)f(x)dx$ (where $f(x)$ is the pdf of X). (The latter, in particular, can be approximated by $n^{-1} \sum_{i=1}^n \hat{r}^2(x_i)$.) Our testing procedure would then amount to (i) finding the residuals by fitting the parametric model, (ii) smoothing the residuals to get \hat{r} , (iii) calculating the size of \hat{r} , and (iv) simulating to get a distribution for how big \hat{r} should be, under the null hypothesis that the parametric model is right.

An alternative to measuring the size of the expected-residuals function would be to try to predict the residuals. We would compare the MSEs of the “model” that the residuals have conditional expectation 0 everywhere, to the MSE of the model that predicts the residuals by smoothing against X , and proceed much as before⁹.

Stabilizing the Sampling Distribution of the Test Statistic

I have just looked at the difference in MSEs. The bootstrap principle being invoked is that the sampling distribution of the test statistic, under the estimated parametric model, should be close to the distribution under the true parameter

⁷ This curse manifests itself here as a loss of power in the test. Said another way, because unconstrained non-parametric regression must use a lot of data points just to determine the general shape of the regression function, even more data is needed to tell whether a particular parametric guess is wrong.

⁸ If you’ve taken functional analysis or measure theory, you may recognize these as the (squared) L_2 and $L_2(f)$ norms of the function \hat{r} .

⁹ Can you write the difference in MSEs for the residuals in terms of either of the measures of the size of \hat{r} ?

value. As discussed in Chapter 6 sometimes some massaging of the test statistic helps bring these distributions closer. Some modifications to consider:

- Divide the MSE difference by an estimate of the noise σ .
- Divide by an estimate of the noise σ times the difference in degrees of freedom, using the effective degrees of freedom (§1.5.3.2) of the nonparametric regression.
- Use the log of the ratio in MSEs instead of the MSE difference.

Doing a double bootstrap can help you assess whether these are necessary.

9.2 Why Use Parametric Models At All?


It might seem by this point that there is little point to using parametric models at all. Either our favorite parametric model is right, or it isn't. If it is right, then a consistent nonparametric estimate will eventually approximate it arbitrarily closely. If the parametric model is wrong, it will not self-correct, but the nonparametric estimate will eventually show us that the parametric model doesn't work. Either way, the parametric model seems superfluous.

There are two things wrong with this line of reasoning — two good reasons to use parametric models.

1. One use of statistical models, like regression models, is to connect scientific theories to data. The theories are ideas about the mechanisms generating the data. Sometimes these ideas are precise enough to tell us what the functional form of the regression should be, or even what the distribution of noise terms should be, but still contain unknown parameters. In this case, the parameters themselves are substantively meaningful and interesting — we *don't* just care about prediction.¹⁰
2. Even if all we care about *is* prediction accuracy, there is still the bias-variance trade-off to consider. Non-parametric smoothers will have larger variance in their predictions, at the same sample size, than correctly-specified parametric models, simply because the former are more flexible. Both models are converging on the true regression function, but the parametric model converges faster, because it searches over a more confined space. In terms of total prediction error, the parametric model's low variance plus vanishing bias beats the non-parametric smoother's larger variance plus vanishing bias. (Remember that this is part of the logic of testing parametric models in the previous section.) In the next section, we will see that this argument can actually be pushed further, to work with not-quite-correctly specified models.

Of course, both of these advantages of parametric models only obtain *if* they are well-specified. If we want to claim those advantages, we need to check the specification.

¹⁰ On the other hand, it is not uncommon for scientists to write down theories positing linear relationships between variables, not because they actually believe that, but because that's the only thing they know how to estimate statistically.



```
h <- function(x) { 0.2 + 0.5*(1+sin(x)/10)*x }
curve(h(x),from=0,to=3)
```

Figure 9.6 Graph of $h(x) = 0.2 + \frac{1}{2} \left(1 + \frac{\sin x}{10}\right) x$ over $[0, 3]$.

9.2.1 Why We Sometimes Want Mis-Specified Parametric Models

Low-dimensional parametric models have potentially high bias (if the real regression curve is very different from what the model posits), but low variance (because there isn't that much to estimate). Non-parametric regression models have low bias (they're flexible) but high variance (they're flexible). If the parametric model is true, it can converge *faster* than the non-parametric one. Even if the parametric model isn't quite true, a small bias plus low variance can sometimes still beat a non-parametric smoother's smaller bias and substantial variance. With enough data the non-parametric smoother will eventually over-take the mis-specified parametric model, but with small samples we might be better off embracing bias.


To illustrate, suppose that the true regression function is

$$\mathbb{E}[Y|X = x] = 0.2 + \frac{1}{2} \left(1 + \frac{\sin x}{10}\right) x \quad (9.6)$$

This is very nearly linear over small ranges — say $x \in [0, 3]$ (Figure 9.6).

I will use the fact that I know the true model here to calculate the actual expected generalization error, by averaging over many samples (Example 23).

Figure 9.7 shows that, out to a fairly substantial sample size (≈ 500), the lower bias of the non-parametric regression is systematically beaten by the lower variance of the linear model — though admittedly not by much.



```

sizes <- c(5, 10, 15, 20, 25, 30, 50, 100, 200, 500, 1000)
generalizations <- sapply(sizes, nearly.linear.generalization)
plot(sizes, sqrt(generalizations[1, ]), type = "l", xlab = "n", ylab = "RMS generalization error",
     log = "xy", ylim = range(sqrt(generalizations)))
lines(sizes, sqrt(generalizations[2, ]), lty = "dashed")
abline(h = 0.15, col = "grey")

```

Figure 9.7 Root-mean-square generalization error for linear model (solid line) and kernel smoother (dashed line), fit to the same sample of the indicated size. The true regression curve is as in [9.6](#) and observations are corrupted by IID Gaussian noise with $\sigma = 0.15$ (grey horizontal line). The cross-over after which the nonparametric regressor has better generalization performance happens shortly before $n = 500$.

```
nearly.linear.out.of.sample = function(n) {
  x <- seq(from = 0, to = 3, length.out = n)
  y <- h(x) + rnorm(n, 0, 0.15)
  data <- data.frame(x = x, y = y)
  y.new <- h(x) + rnorm(n, 0, 0.15)
  sim.lm <- lm(y ~ x, data = data)
  lm.mse <- mean((fitted(sim.lm) - y.new)^2)
  sim.np.bw <- npregbw(y ~ x, data = data)
  sim.np <- npreg(sim.np.bw)
  np.mse <- mean((fitted(sim.np) - y.new)^2)
  mses <- c(lm.mse, np.mse)
  return(mses)
}

nearly.linear.generalization <- function(n, m = 100) {
  raw <- replicate(m, nearly.linear.out.of.sample(n))
  reduced <- rowMeans(raw)
  return(reduced)
}
```

CODE EXAMPLE 23: *Evaluating the out-of-sample error for the nearly-linear problem as a function of n , and evaluating the generalization error by averaging over many samples.*

9.3 Further Reading

This chapter has been on specification testing for regression models, focusing on whether they are correctly specified for the conditional expectation function. I am not aware of any other treatment of this topic at this level, other than the not-wholly-independent [Spain *et al.* \(2012\)](#). If you have somewhat more statistical theory than this book demands, there are very good treatments of related tests in [Li and Racine \(2007\)](#), and of tests based on smoothing residuals in [Hart \(1997\)](#).

Econometrics seems to have more of a tradition of formal specification testing than many other branches of statistics. [Godfrey \(1988\)](#) reviews tests based on looking for parametric extensions of the model, i.e., refinements of the idea of testing whether $\theta_3 = 0$ in Eq. [9.3](#). [White \(1994\)](#) combines a detailed theory of specification testing within parametric stochastic models, *not* presuming any particular parametric model is correct, with an analysis of when we can and cannot still draw useful inferences from estimates within a mis-specified model. Because of its generality, it, too, is at a higher theoretical level than this book, but is strongly recommend. White was also the co-author of a paper [\(Hong and White 1995\)](#) presenting a theoretical analysis of the difference-in-MSEs test used in this chapter, albeit for a particular sort of nonparametric regression we've not really touched on.

Appendix [F](#) considers some ways of doing specification test for models of distributions, rather than regressions.