10 (Almost) everything you need to know about Curves & Surfaces

Curves arise in nearly every aspect of securities and derivatives trading. In many cases the curves don't really exist but one may imagine that a series of points form a curve that results from "connecting the dots". This is an essential requirement for most valuation and risk management issues where interpolation is required. This Chapter presents these ideas, with emphasis on the notion that that such curves are actually "models" of our view of the world. As always, blindly applying such methods can be dangerous, and so there are suggested guidelines and some special methods for dealing with the real world.

This Chapter introduces the two primary methodologies for fitting curves and surfaces: interpolation methods and approximation methods. These represent two different classes of methods. The major distinction between these two classes[3] may be summarised as:

- Interpolation methods: fitting curves exactly[4] through the data points (e.g. splines).
- Approximation methods: fitting the "trend" in curves (e.g. least squares)

Each class aims to capture "information" in the data. However, each class of methods is appropriate for different types of data (e.g. noisy data from market histories vs. exact data from benchmark bond rates for yield curves), or to meet different objectives (interpolation between two points vs. extrapolating trends)

There is also a distinction between the "how's and why's" of the "physics" in the data. That is, in some cases one is simply trying to extract hitherto unknown characteristics from the dataset. In other cases, one may reasonably expect a particular type of relationship, and so there is an apriori model that needs to be confirmed, or parameterised. This distinction will sometimes be referred to as "empirical" fitting vs. "fundamental" fitting[5].

---

[3] A very important class of problems related to the representation of curves and surfaces that act as approximating functions as building block for more complex techniques. Notably, the creation of Finite Element and Finite Difference representations will rely on a range of simply to potentially quite complex curve/surface fitting technologies to approximate solutions to Partial Differential Equations.

[4] Although this may sound a bit like series expansion, it is a fundamentally different process. In particular, interpolation methods only aim (at most) to fit a curve that is exact at the data. That is crucially a different proposition from the objectives, where the aim is to recreate the entire function (i.e. at every point and for every derivative) exactly.

[5] Time Series Analysis is a good example of how this type of distinction may arise, where Econometrics is a word used to "fit" "fundamental" models of economic dynamics (e.g. GDP as function of employment, inflation, etc), while traditional time series estimates of parameters for "guessed models" are an empirical exercise (e.g. tomorrow's volatility forecast is guessed to some linear model and the parameters are extracted empirically).

Examples of application of interpolation methods in trading and risk management include:

- Yield curve generation and interpolation
- Volatility skew generation and interpolation
- Volatility surfaces for term-structure options

Notice the potentially confusing terminology; an "interpolation methods" is class of techniques, while "interpolating" (e.g. between two points) is the application of an interpolation method.

Examples of application of approximation methods in trading and risk management include:

- Cheap/dear analysis
- Trend analysis
- Volatility cones

The creation and assessment (e.g. goodness of fit) of these methods relies heavily on the "basic maths & stats" in Part I of this Book. Importantly, while the series representation/expansion discussions in Chapter 8 at first may appear to be a curve fitting matter, they are considerably more. This Chapter is only concerned with the considerably simpler matter of producing a curve that "looks close enough" to the data. Chapter 8, on the other hand, is concerned with an exact functional representation.

10.1 Curve & Surface Fitting: The Basic Idea

The basic idea for all curve and surface fitting is to relate a functional form to a dataset. Technically, given a set of data $\{d_1, d_2, d_3, \ldots d_n\}$ (possibly vector valued), find a function and parameters which "best represents" the information in the data. It is the specific definition or choice of "best represents" that differentiates the methodologies.

The "big picture" issues associated with curve fitting, and the primary factors differentiating interpolation methods and approximation are provided here.

10.1.1 Graphical Illustration of Interpolation and Approximation

Figure 10.1 – 1 a) and b) illustrate two different curve fitting problems. Figure 10.1 – 1 a) shows data that is either “noisy” or “stochastic” (e.g. a first return map as seen in Section 5.1.2), and a curve has been generated that captures the trend in the data. Figure 10.1 – 1 b) show a data set that is deterministic and may well be expected to have a known “shape” (e.g. a yield curve), and so the curve has been fit that passes through each point.

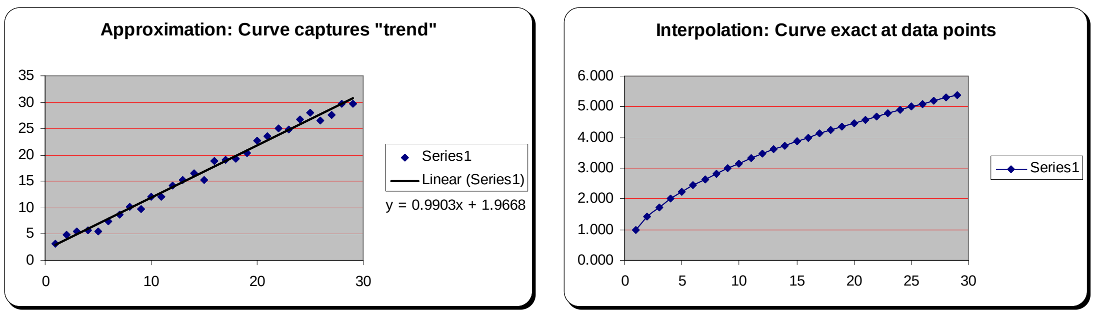

Figure 10.1 – 1 a) Approximation of “noisy” data, b) Interpolation of “exact” data.

This illustrates the primary difference between the two classes of curve fitting methodologies.

Other “big picture” issues are worthy of mention in regards to the differences and idiosyncrasies between approximation method and interpolation methods.

With approximation methods, one major issue is the how to decide between two competing “fits” that may have equally good “explanatory powers”. For example, Figure 10.1 – 2 a) and b) illustrates two possible curves fit through the same noisy dataset. One fitting curve is linear, while the other fitting curve is quadratic. It can happen that the mathematical tools for deciding “goodness of fit” indicate the same or very similar goodness measures for two different curves. In those cases, other tools are required to “decide”. Importantly, there may be no way to prove if one or the other of the “guessed” curves is “correct” (at least not without additional information, which may or may not exist).

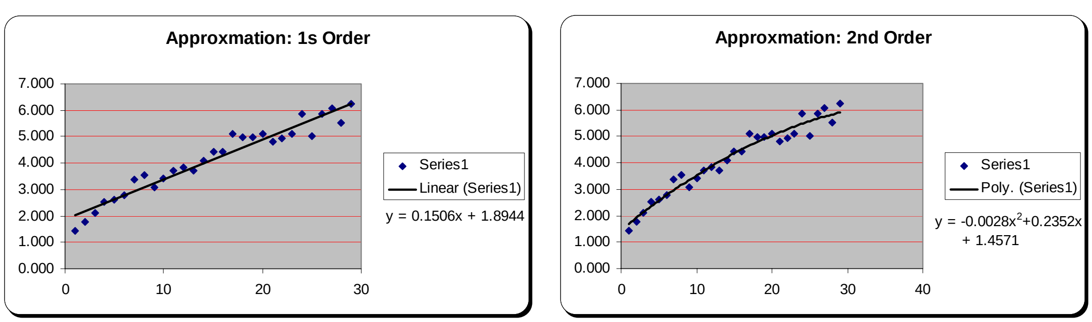

Figure 10.1 – 2 a) Linear approximation of “noisy” data, b) Quadratic approximation of “noisy” data.