NON-LINEAR SYSTEMS

- **Topic:** How to efficiently (and accurately) solve a systems of non-linear equations.
  - Almost all static economic models can be characterized as systems of nonlinear algebraic equations.
  - Moreover, the solution of nonlinear equations is often an essential intermediate step in other procedures.
- Hence, we will now consider the problem of finding a zero of a system of nonlinear equations:

$$F(x) \equiv \begin{bmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_n(x) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad x \in X$$

Preliminaries: Banach's Theorem

Definition *Let X be a normed vector space. An operator* $T : D \subseteq X \to X$ *is a contraction operator (or contraction mapping) if there exists a* $\gamma \in [0,1)$ *such that* $\| T(x_1) - T(x_2) \| \leq \gamma \| x_1 - x_2 \|$ *for all* $x_j \in D$.

Definition *An element* $\hat{x} \in X$ *is a fixed point for an operator* $T : D \subseteq X \to X$ *if* $T(\hat{x}) = \hat{x}$.

Theorem (Banach) *If X is a normed vector space, D a complete subset of X, and* $T : D \to D$ *a contraction operator, then T has a unique fixed point.*

Fixed point iteration

- The proof of Banach's Theorem is the theoretical basis of the **fixed point iteration** (or **successive approximations**) solution method for fixed point problems:

Corollary *Let D be a complete subset of a normed vector space X, and* $T : D \to D$ *a contraction operator. The successive approximations:*

$$x_{k+1} = T(x_k), \quad k = 0, 1, 2, \ldots$$

converge to the unique fixed point of T for any initial guess $x_0 \in D$*.*

- Some general sufficient conditions for a contraction are available in the literature (see Stokey and Lucas 1989, Th. 3.3, p. 54).

Convergence and stopping rules

- Successive approximation schemes convergence only asymptotically.
- Any iterative algorithm therefore needs a feasible **stopping rule**, i.e. a rule that terminates the iteration when a sufficiently good approximation has been reached.
- The most useful general stopping rule requires the iteration to stop as soon as the percentage change in $||x||$ becomes small relatively to some tolerance parameter ε.
- In other words, the iteration stops and returns the result as soon as $||\Delta x_{k+1}|| \leq \varepsilon(1+||x_k||)$, where the unit in the right-hand side takes care of the possibility that x goes to zero.

This rule is essentially based on the following result:

Lemma *Let $x_k \in R^n$ for $k = 0, 1, 2, \ldots, \infty$. If the sequence $\{x_k\}_{k=0}^{\infty} \in R^{\infty}$ converges superlinearly to $\hat{x} \in R^n$, then:*

$$\lim_{k \to \infty} \frac{\|x_{k+1} - x_k\|}{\|x_k - \hat{x}\|} = 1$$

for any norm in R^n.

Remark *In other words, if a sequence of real vectors converges at least superlinearly, then in the limit the size of the step, $\|x_{k+1} - x_k\|$, is essentially equal to the size of the approximation error, $\|x_k - \hat{x}\|$.*

Fixed point iterations for non-linear equations

- Our problem can easily be transformed into a fixed point problem; define:

$$G(x) \equiv F(x) + x = x$$

- If X is complete and G a contraction operator on X, Banach's Theorem guarantees that the successive approximation scheme:

$$x_{k+1} = G(x_k), \quad k = 0, 1, 2, \ldots$$

converges to the unique solution for any $x_0 \in X$.

- Of course, the point is that G is not necessarily a contraction operator.

- We present now a sufficient condition for a contraction, based on the assumption that G is a C^1 operator:

Theorem (Mean Value) *Let X be an open and convex subset of R^n, and let $G : X \to R^n$ be a C^1 operator. Then:*

$$\| G(x) - G(y) \| \leq \max_{\lambda \in [0,1]} \| J_G[\lambda x + (1-\lambda) y] \| \, \| x - y \|$$

for all $x, y \in X$ and all norms on R^n.

Theorem *Let X be a bounded, closed, convex, and nonempty subset of R^n, and let $G : X \to X$ be a C^1 operator. If:*

$$\gamma \equiv \max_{x \in X} \| J_G(x) \| < 1$$

for some norm on R^n, then G is a (differentiable) contraction operator with Lipschitz constant γ.

Theorem (Ostrowski) *Let $\hat{x} \in X$ be the fixed point of a C^1 operator $G : X \to X$ such that $\| J_G(\hat{x}) \| < 1$ for some norm on R^n. Then the successive approximation scheme is locally convergent.*

Newton's method

- Newton's method proceeds by **successive linearizations**:
    - at each iteration, the original system is linearized around the current guess x_k, and the linear system is typically solved using Gaussian elimination;
    - the result is then used as the initial guess for the next iteration.
- Under some conditions, the method is locally **quadratically convergent**, and therefore convergence can be assessed with a standard stopping rule.

- More formally, let X be an open subset of R^n, and assume that $F:X\rightarrow R^n$ is a C^1 operator.
- Furthermore, assume that $J(x)$, the Jacobian of F evaluated at x, is nonsingular for all x in X
- Given an initial guess x_0 in X, the first-order Taylor expansion around x_0 is

$$F(x)\approx F(x)\equiv F(x_0)+J(x_0)(x-x_0)$$

- The linear problem $F(x)=0$ can be rewritten as

$$J(x_0)(x_0-x)=F(x_0)$$

and solved using Gaussian elimination.

- The resulting x_1 can then be used as the initial guess for a new iteration.

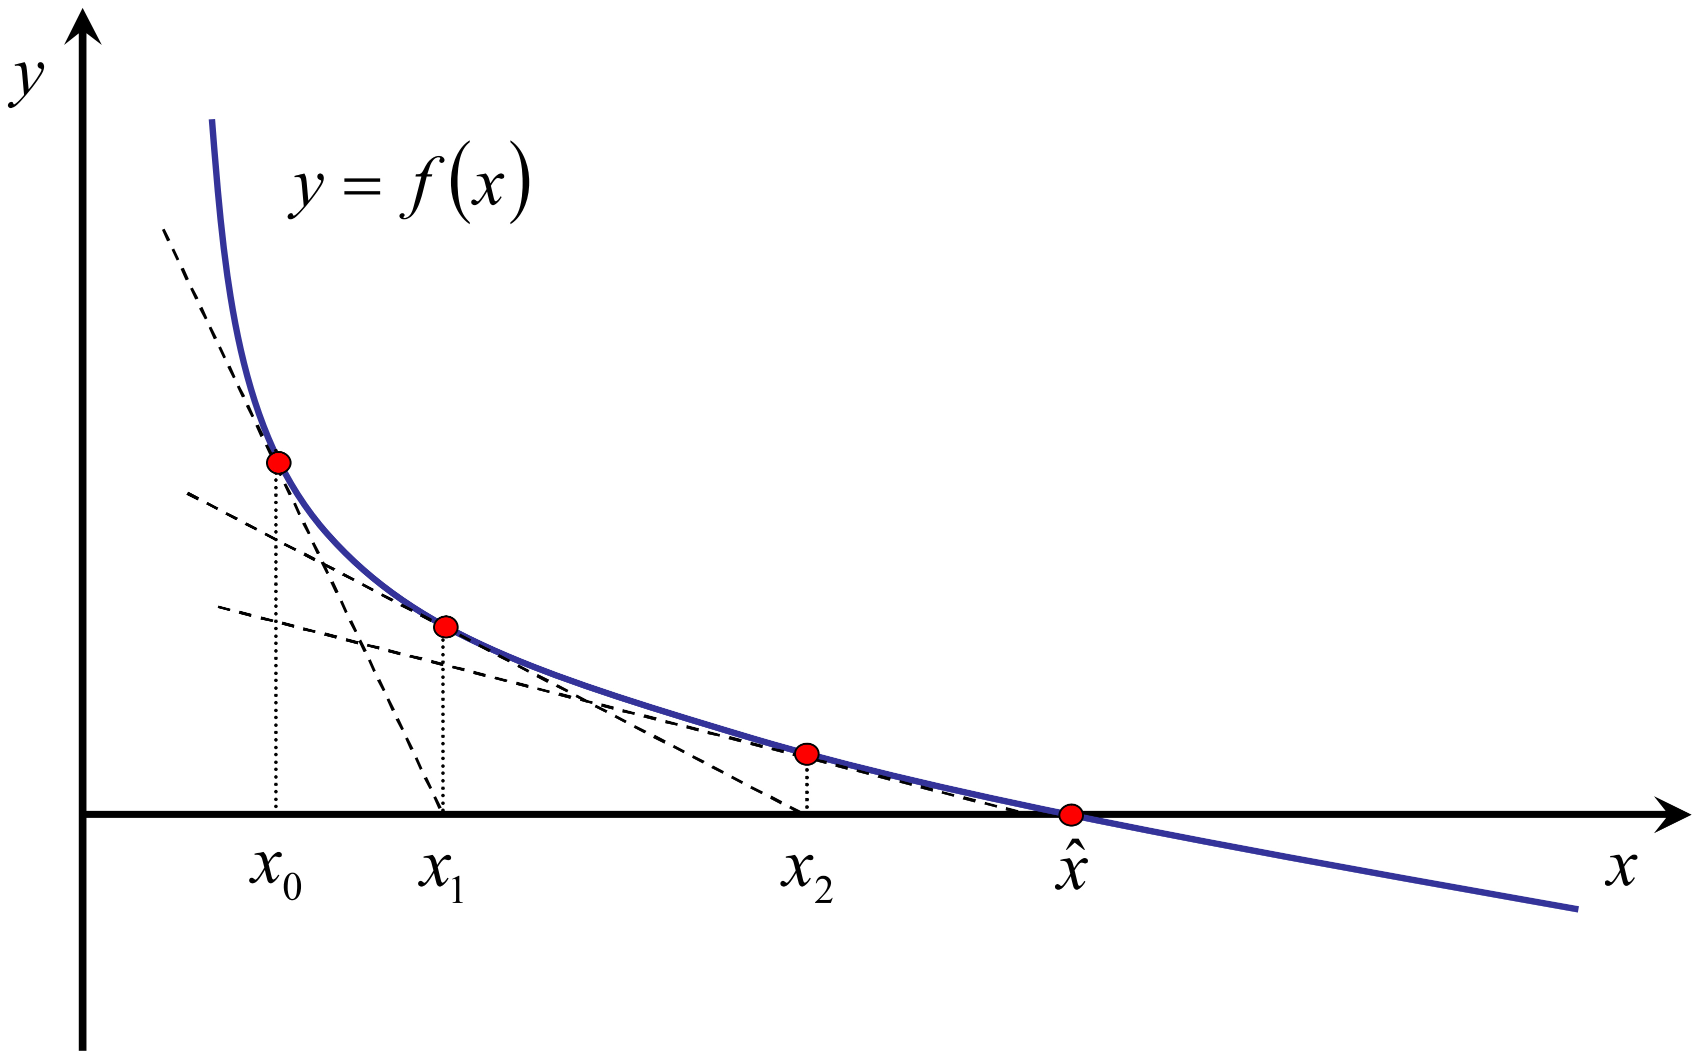
y
$y = f(x)$
x_0
x_1
x_2
$\hat{x}$
x

- Hence, Newton's method can be described by the following successive approximation scheme:

$$x_{k+1} = G(x_k) \equiv x_k - J(x_k)^{-1} F(x_k), \qquad k=0,1,2,\ldots$$

- The vector:

$$d_k \equiv \Delta x_{k+1} = - J(x_k)^{-1} F(x_k)$$

is known as **Newton's step**.

Theorem *Let* $F : X \subseteq R^n \to R^n$ *be a* C^1 *operator. If:*

i) the equation $F(x) = 0$ *has a unique solution* $\hat{x} \in X$*;*

ii) $J : X \to R^{n \times n}$ *is a Lipschitz continuos operator with Lipschitz constant* $\gamma > 0$*;*

iii) $J(\hat{x})$ *is nonsingular;*

then there is a $\zeta > 0$ *such that* $G : B(\zeta) \to B(\zeta)$*, where* $B(\zeta) \equiv \{x : \|x - \hat{x}\| < \zeta\} \in X$ *and* $G(x_k) \equiv x_k - J_k^{-1} F(x_k)$ *is a contraction operator.*

Corollary *Assume that the hypotheses of the previous Theorem hold. Newton's successive approximation scheme converges quadratically:*

$$\|x_{k+1} - \hat{x}\| \leq \varphi \|x_k - \hat{x}\|^2, \quad k = 0, 1, 2, \ldots$$

where $\varphi \equiv \gamma \|J(\hat{x})^{-1}\|$*, for all* $x_0 \in B(\zeta)$*.*

- Given the quadratic convergence of Newton's method, the general stopping rule:

$$|| d_k || \leq \varepsilon \, (1 + || x_k ||)$$

where $\varepsilon>0$, will stop the iterations as soon as the approximation error $|| x_k - x ||$ is of order ε, unless the system is particularly ill-behaved.

- However, if the initial guess is not good enough, i.e. if $x_0 \notin B$, Newton's method may fail to converge to a zero of F.

- Since the size of B is generally unknown ex-ante, we should consider x_k a solution only if:

$$|| F(x_k) || \leq \delta \, (1 + || F (x_0) ||)$$

where $\delta>0$ is another tolerance parameter.

Finite differences

- A critical step in Newton's method requires the computation of the Jacobian matrix of F at a given x.
- Often the Jacobian can not be easily computed analytically: in these cases, a numerical approach is needed.
- Numerical differentiation is an essential application of the **finite difference method**.

- Assume that f:$X \rightarrow R$ is C^k on X
- Consider the one-sided Taylor expansion of $f(x)$ around an arbitrary point x in X (where $h>0$):

$$f(x+h) = f(x) + hf'(x) + \frac{1}{2}h^2 f''(x) + \frac{1}{6}h^3 f'''(x) + \ldots$$

- The previous expression can be rewritten as:

$$\boxed{f'(x) = \frac{f(x+h) - f(x)}{h} + O\left(\frac{hf''(x)}{2}\right)}$$

- This expression is known as the one-sided finite difference formula: the $O[(h/2)f''(x)]$ term on right-hand side is called **truncation error**.

- Truncation is not the only source of error: the round-off error is, as always, an additional source of inaccuracy.

- The round-off error is on the order of $\varepsilon_f \,|\, f(x)/h \,|$, where ε_f is the accuracy with which f is computed: generally, it is comparable to ε_M, the machine's internal precision.

- The truncation error, instead, is on the order of $|(h/2)f''(x)|$.

- These two errors can be jointly minimized by choosing:

$$h = \sqrt{\varepsilon_M \left| \frac{2f(x)}{f''(x)} \right|} = \sqrt{\varepsilon_M}\, x_c$$

where $x_c \equiv (\,|\, 2\, f(x) \,/\, f''(x) \,|\,)^{½}$.

- If no specific information on the curvature of f is available, a standard choice is $x_c = 1 + |x|$.

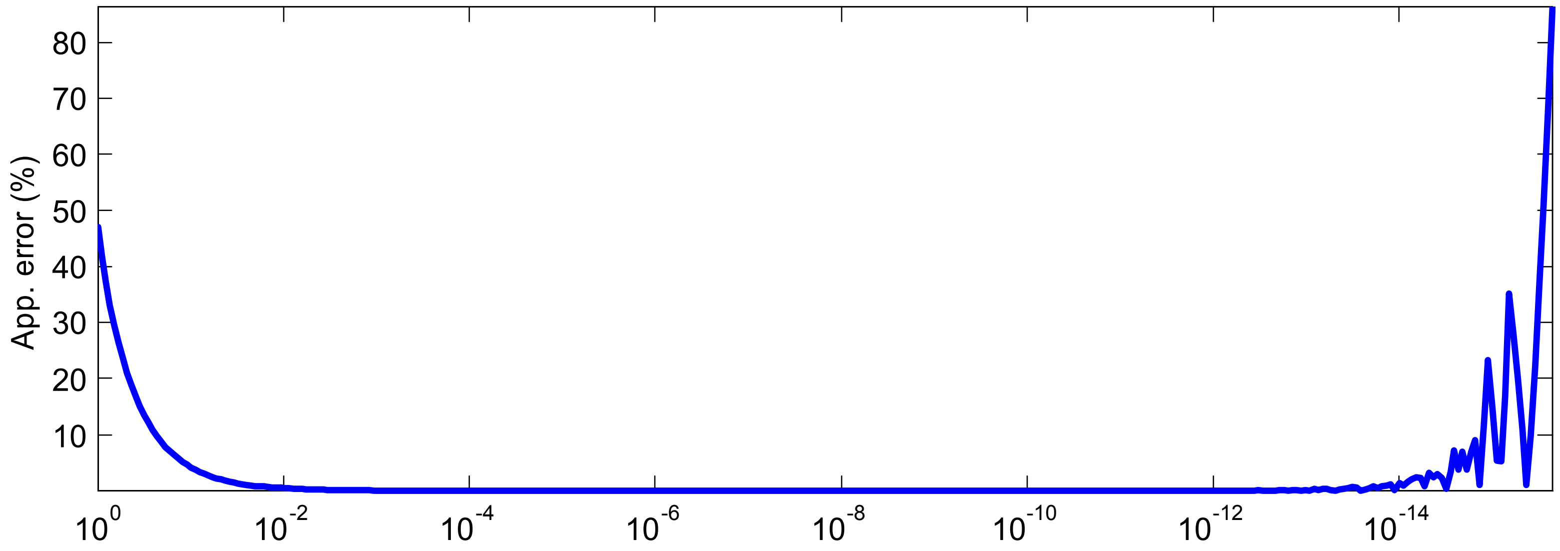

App. error (%)
80
70
60
50
40
30
20
10
10^0
10^-2
10^-4
10^-6
10^-8
10^-10
10^-12
10^-14

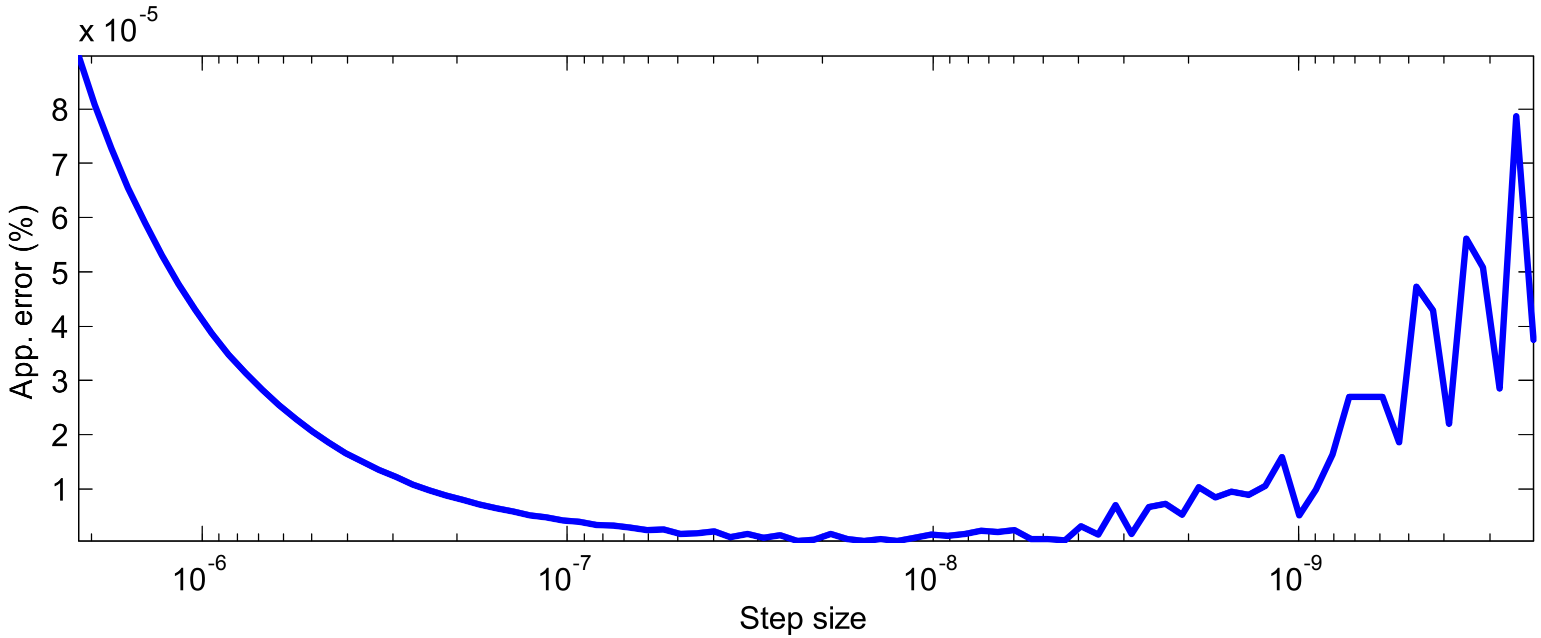

x 10^-5
App. error (%)
8
7
6
5
4
3
2
1
10^-6
10^-7
10^-8
10^-9
Step size

- We can combine the previous one-sided Taylor expansion with the symmetric one to points to the left of x to obtain:

$$f(x+h) - f(x-h) = 2hf'(x) + \frac{1}{3}h^3 f'''(x) + \ldots$$

$$\boxed{f'(x) = \frac{f(x+h) - f(x-h)}{2h} + O\left(\frac{h^2 f'''(x)}{3}\right)}$$

- This is known as the **two-sided finite difference formula**.
- Note that the truncation error is in this case on the order of h^2 instead of h.
- The truncation and round-off errors are minimized by choosing:

$$h = \sqrt[3]{\varepsilon_M \left| \frac{3f(x)}{f'''(x)} \right|}$$

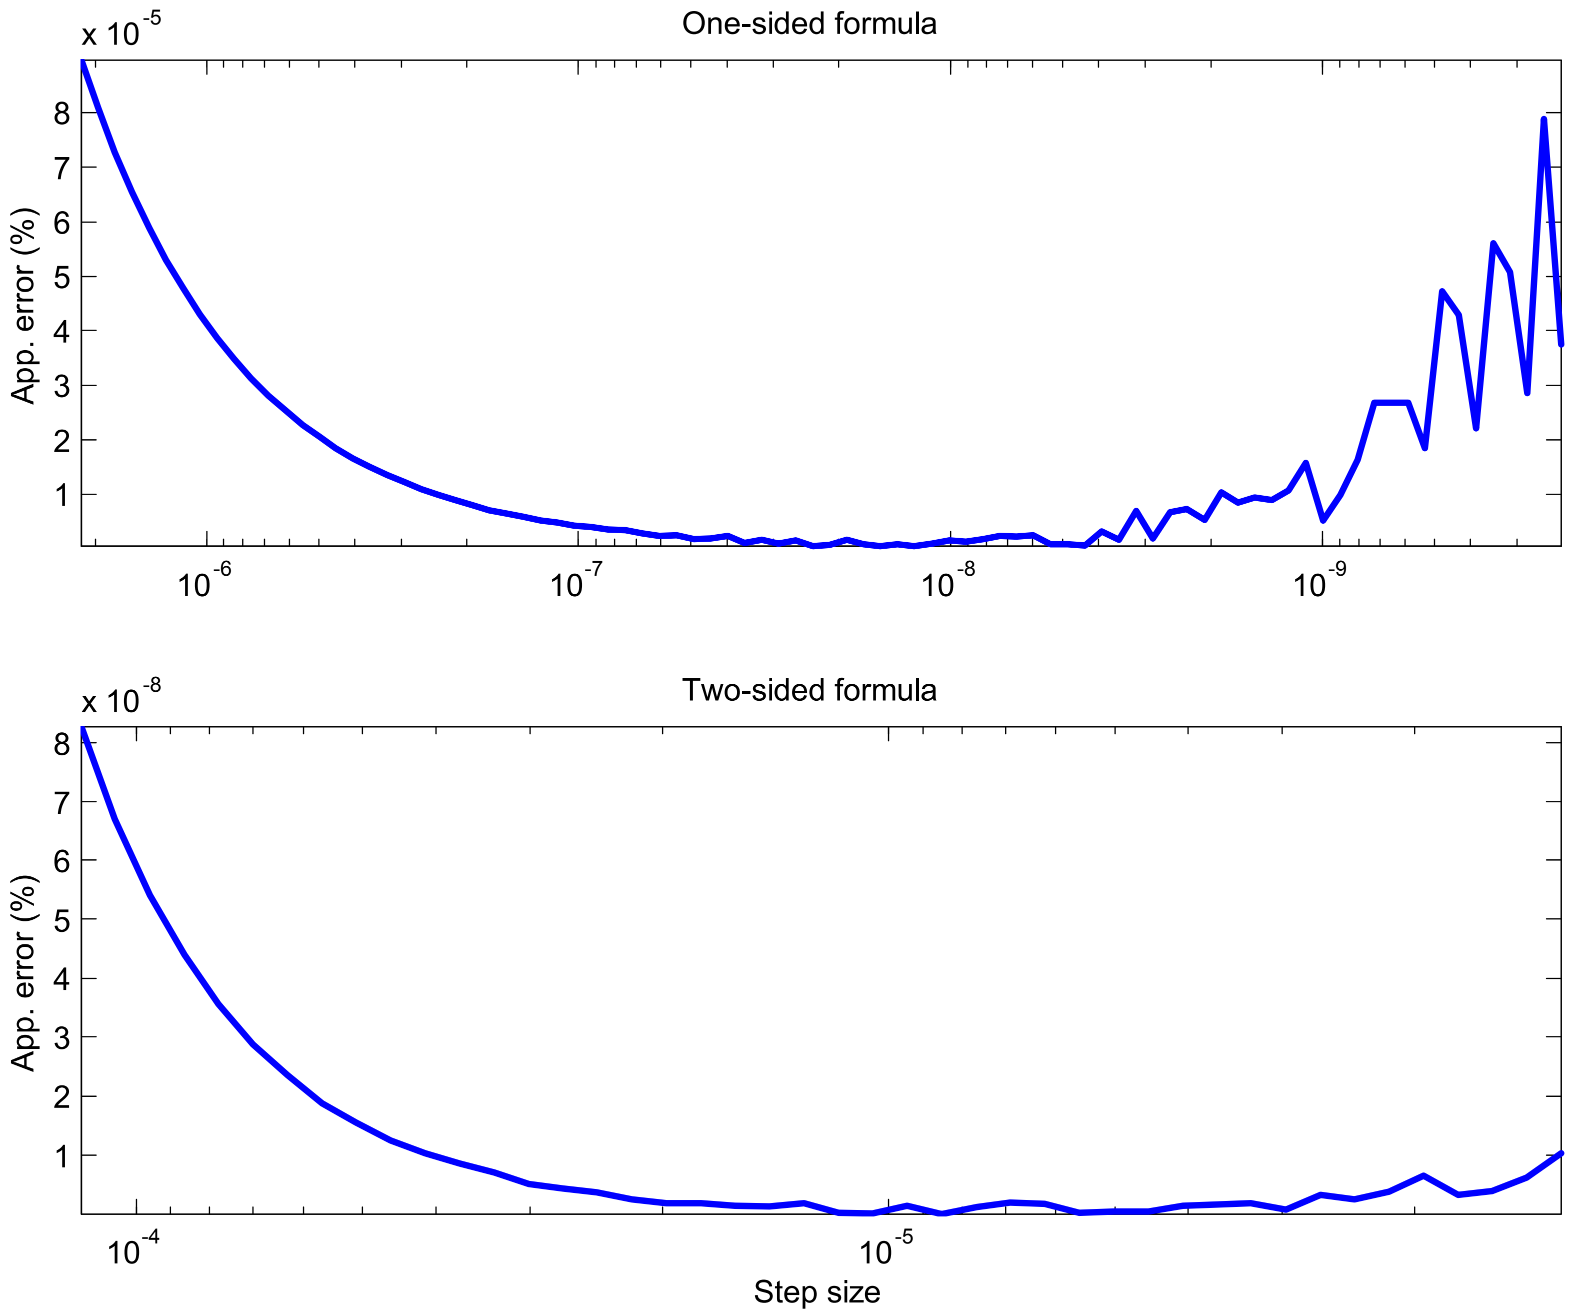
One-sided formula
x 10^-5
App. error (%)
8
7
6
5
4
3
2
1
10^-6
10^-7
10^-8
10^-9
Two-sided formula
x 10^-8
App. error (%)
8
7
6
5
4
3
2
1
10^-4
10^-5
Step size

- Assume now that $F: R^n \rightarrow R^m$ is C^k at x
- The Jacobian of F can be numerically computed using a one-sided finite difference formula:

$$\boxed{\frac{\partial f_i(x)}{\partial x_j} = \frac{f_i(x + h_j e_j) - f_i(x)}{h_j} + O(h_j)}$$

where $i=1,2,\ldots,m$, $j=1,2,\ldots,n$, e_j is a column vector of zeros with just its j_{th} element equal to one, and:

$$h_j = \sqrt{\varepsilon_F}\,(1 + |x_j|)$$

Quasi-Newton Methods

- Numerical differentiation is the most computationally expensive step in Newton's method.
- Quasi-Newton methods use approximations of the Jacobian, gaining in computational efficiency but loosing the quadratic convergence of Newton’s method.

Frozen Newton's method

- The simplest Quasi-Newton method is the so-called frozen (or simplified) Newton's method:

$$d_k = -J_0^{-1} f(x_k), \quad k = 0, 1, 2, \ldots$$

- This scheme is quite unstable and converges only linearly.

Broyden's Method

- Assume that f:$R \rightarrow R$. Let x and y be two points in R.
- The first order derivative of f could be roughly approximated near x and y by the slope of the secant line:

$$f' \approx \hat{f}' \equiv \frac{f(x) - f(y)}{x - y}$$

- We could simplify Newton's method by substituting the actual derivative with the equivalent of f':

$$d_k = -\frac{\Delta x_k}{f(x_k) - f(x_{k-1})} f(x_k), \quad k = 1, 2, \ldots$$

- This *secant method* needs an initial condition for the derivative: the obvious choice is a numerically computed derivative at x_0.

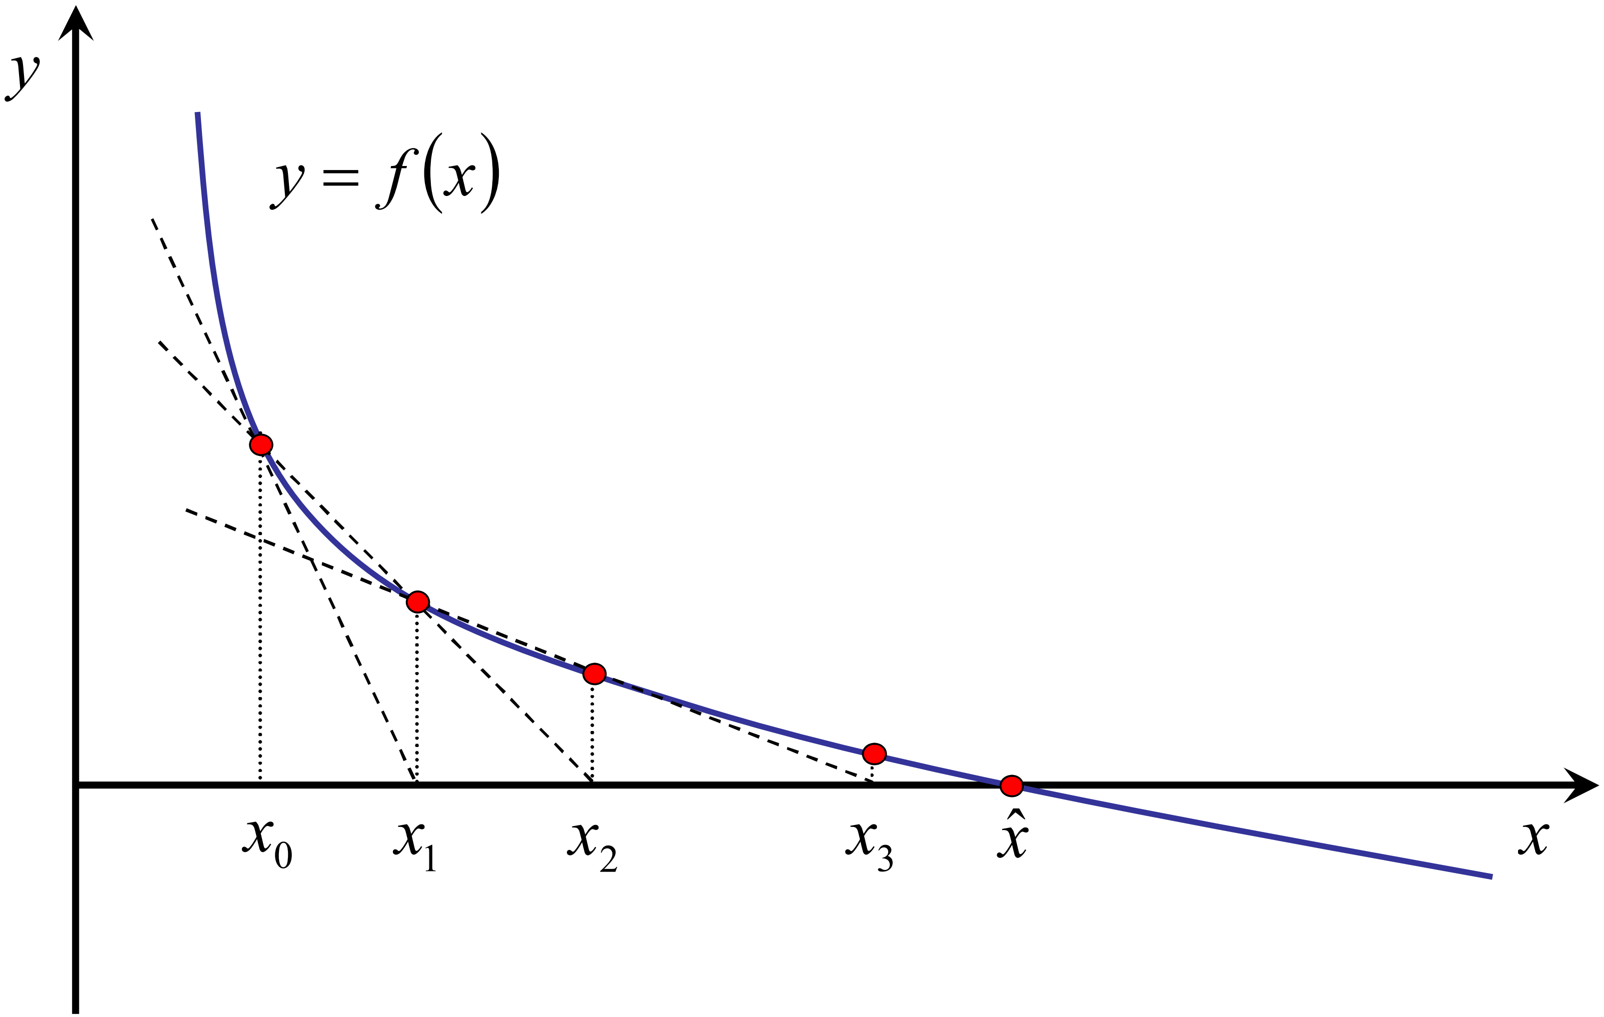

y
$y = f(x)$
x_0
x_1
x_2
x_3
$\hat{x}$
x

- The secant method cannot be directly extended to the multivariate case.
- Assume that $F\colon R^n \to R^n$. Let x and y be two points in R^n and $J(z)$ the Jacobian of F at z.
- We can show that the Jacobian approximately solves, near x (or y):

$$F(x) - F(y) \approx J(x)(x - y)$$

- However, given x, y, $F(x)$, and $F(y)$ this *secant equation* is not enough to pin down an approximation of $J(x)$ if $n>1$:
  - $F(x)-F(y)$ and $x-y$ are **column vectors**, and therefore the equation imposes only n constraints, while $J(x)$ has n^2 unknown.

Given an initial guess for the Jacobian at x_0, J_0, Broyden's method iterates on the following successive approximation scheme:

$$d_k = -J_k^{-1}F(x_k), \quad k = 0, 1, 2, \ldots$$

where d_k is generically known as *quasi-Newton step*.

The approximated Jacobian is updated at each iteration by imposing two conditions:

1. the update J_{k+1} has to be a good approximation of the Jacobian near x_k and x_{k+1}; hence, it has to solve the secant equation:

$$F(x_{k+1}) - F(x_k) = J_{k+1}d_k$$

2. the change between J_{k+1} and J_k has to be the smallest possible according to the *Frobenius matrix norm*, defined as:

$$\|X\|_F \equiv \sqrt{\sum_{j,i=1}^{n} x_{ji}^2}$$

- The second requirement is based on the observation that the secant equation is the only new piece of information that becomes available at each iteration.
- The iteration scheme should preserve as much as possible of the information already acquired (summarized by the current approximation J_k)
- The Frobenius norm takes into account changes to all elements of J.
- The corresponding updating rule is:

$$\Delta J_{k+1} = \frac{F(x_{k+1})d_k'}{d_k' d_k}$$

- Broyden's method is much less computationally intensive than Newton's method.
- However, this implies also slower convergence: Broyden's method can be shown to converge superlinearly when the initial approximation of the Jacobin is good enough.
- Thanks to the *Sherman-Morrison-Woodbury Lemma*, we can obtain an updating rule for the inverse of the Jacobian:

$$\Delta J_{k+1}^{-1} = \frac{z_k d_k' J_k^{-1}}{d_k'(d_k - z_k)}$$

where:

$$z_k \equiv -J_k^{-1} F(x_{k+1})$$