Introduction to Longitudinal Analysis

James H. Steiger

Department of Psychology and Human Development
Vanderbilt University

GCM, 2010

A Conceptual Framework

The study of change is central to educational and developmental psychology. We ask questions like:

1. How do certain skills develop?
2. At what rate do skills change?
3. When are they most likely to change?
4. What factors influence change?

A Conceptual Framework

The study of change is also very important in abnormal psychology. We ask questions like:

1. What is the typical course of drug and alcohol use among teenagers? Is change over time linear or nonlinear?
2. What factors influence smoking behavior during adolescence and young adulthood?
3. What is the typical developmental sequence of adolescent depression? How does it vary between boys and girls?

A Conceptual Framework

In virtually any area of psychology or education, we sooner or later find ourselves in a position where a coherent and accurate approach to the study of change is vital.

Some Key Conceptual Distinctions

In studying change, we shall be concerned with some key conceptual distinctions:

1. *Individual vs. Population.*We are interested in the overall pattern of change manifested by our population(s) of interest. However, we are also interested in how individual trajectories vary, and why they vary.
2. *Trajectory vs. Covariate.* After characterizing the trajectories of our population, we seek covariates that reliably predict the characteristics of those trajectories, and the precise functional nature of how the covariates predict.

Changes in Antisocial Behavior During Adolescence

- Adolescence is an eventful time in the lives of many people
- Most emerge with a few scars, but basically healthy
- A minority of teenagers exhibit antisocial behaviors, including depressive *internalizing* and hostile, aggressive *externalizing* behaviors
- Recent advances in statistical methods have led to empirical exploration of developmental trajectories, and investigation of variables that predict antisocial behaviors based on early childhood events and symptoms

Changes in Antisocial Behavior During Adolescence

Coie, Terry, Lenox, Lochman, and Hyman (1995, *Development and Psychopathology*, 697–713) studied 407 public school students in Durham, NC.

- Each student was assessed in third grade with screening instruments designed to measure aggressive behavior
- A stratified random sample was selected based on screening results
- In 6th, 8th, and 10th grades, these students were administered a battery of tests, including the CAS (Child Assessment Schedule), which assesses antisocial behaviors.
- Patterns of change, and their predictability from 3rd grade ratings, were studied in parallel analyses of boys and girls.

Changes in Antisocial Behavior During Adolescence

Coie, Terry, Lenox, Lochman, and Hyman (1995, *Development and Psychopathology*, 697–713) studied 407 public school students in Durham, NC.

- Each student was assessed in third grade with screening instruments designed to measure aggressive behavior
- A stratified random sample was selected based on screening results
- In 6th, 8th, and 10th grades, these students were administered a battery of tests, including the CAS (Child Assessment Schedule), which assesses antisocial behaviors.
- Patterns of change, and their predictability from 3rd grade ratings, were studied in parallel analyses of boys and girls.

Changes in Antisocial Behavior During Adolescence

Coie, Terry, Lenox, Lochman, and Hyman (1995, *Development and Psychopathology*, 697–713) studied 407 public school students in Durham, NC.

- Each student was assessed in third grade with screening instruments designed to measure aggressive behavior
- A stratified random sample was selected based on screening results
- In 6th, 8th, and 10th grades, these students were administered a battery of tests, including the CAS (Child Assessment Schedule), which assesses antisocial behaviors.
- Patterns of change, and their predictability from 3rd grade ratings, were studied in parallel analyses of boys and girls.

Changes in Antisocial Behavior During Adolescence

Coie, Terry, Lenox, Lochman, and Hyman (1995, *Development and Psychopathology*, 697–713) studied 407 public school students in Durham, NC.

- Each student was assessed in third grade with screening instruments designed to measure aggressive behavior
- A stratified random sample was selected based on screening results
- In 6th, 8th, and 10th grades, these students were administered a battery of tests, including the CAS (Child Assessment Schedule), which assesses antisocial behaviors.
- Patterns of change, and their predictability from 3rd grade ratings, were studied in parallel analyses of boys and girls.

Changes in Antisocial Behavior During Adolescence

Findings (Boys Data)

Outcomes differed depending on 3rd grade ratings:

- Nonaggressive 3rd grade boys showed essentially no increase in aggressive behaviors between 6th and 10th grades.
- Aggressive nonrejected boys showed a temporary increase in internalizing behaviors, but this declined back to the level of nonaggressive boys by 10th grade.
- Aggressive rejected boys showed a linear increase in both internalizing and externalizing behaviors between 6th and 8th grades.

Changes in Antisocial Behavior During Adolescence

Findings (Boys Data)

Outcomes differed depending on 3rd grade ratings:

- Nonaggressive 3rd grade boys showed essentially no increase in aggressive behaviors between 6th and 10th grades.
- Aggressive nonrejected boys showed a temporary increase in internalizing behaviors, but this declined back to the level of nonaggressive boys by 10th grade.
- Aggressive rejected boys showed a linear increase in both internalizing and externalizing behaviors between 6th and 8th grades.

Changes in Antisocial Behavior During Adolescence

Findings (Boys Data)

Outcomes differed depending on 3rd grade ratings:

- Nonaggressive 3rd grade boys showed essentially no increase in aggressive behaviors between 6th and 10th grades.
- Aggressive nonrejected boys showed a temporary increase in internalizing behaviors, but this declined back to the level of nonaggressive boys by 10th grade.
- Aggressive rejected boys showed a linear increase in both internalizing and externalizing behaviors between 6th and 8th grades.

Testing Hypotheses about Reading Development

Experts are divided on the why children differ markedly in learning to read. Francis, Shaywitz, Stuebing, Shaywitz, and Fletcher (1996, *Journal of Educational Psychology*, 3–17) used multilevel modeling to examine hypotheses about reading development.

Two major competing hypotheses are:

1. The *lag* hypothesis. This assumes that every child can become a proficient reader, and that "children who differ in reading ability vary only in the rate at which cognitive skills develop, so the skill will emerge over time." (Francis, et al., p.3)
2. The *deficit* hypothesis. This hypothesizes "children fail to read proficiently because of the absence of a skill that never develops sufficiently." (Francis, et al., p. 3)

Testing Hypotheses about Reading Development

Experts are divided on the why children differ markedly in learning to read. Francis, Shaywitz, Stuebing, Shaywitz, and Fletcher (1996, *Journal of Educational Psychology*, 3–17) used multilevel modeling to examine hypotheses about reading development.

Two major competing hypotheses are:

1. The *lag* hypothesis. This assumes that every child can become a proficient reader, and that "children who differ in reading ability vary only in the rate at which cognitive skills develop, so the skill will emerge over time." (Francis, et al., p.3)
2. The *deficit* hypothesis. This hypothesizes "children fail to read proficiently because of the absence of a skill that never develops sufficiently." (Francis, et al., p. 3)

Testing Hypotheses about Reading Development

FSSSF(1996) gathered data annually on 363 six-year-olds until they reached the age of 16.

Besides completing the Woodcock-Johnson Psycho-educational Test Battery each year, they also completed the WISC every second year.

Three groups were identified:

1. 301 *normal readers.* Performance was within limits concomitant with WISC score
2. 28 *discrepant readers.* Performance was significantly lower than predicted by WISC score
3. 34 *low achievers.* Performance was low, but not markedly lower than predicted from WISC

Testing Hypotheses about Reading Development

FSSSF(1996) gathered data annually on 363 six-year-olds until they reached the age of 16.

Besides completing the Woodcock-Johnson Psycho-educational Test Battery each year, they also completed the WISC every second year.

Three groups were identified:

1. 301 *normal readers.* Performance was within limits concomitant with WISC score
2. 28 *discrepant readers.* Performance was significantly lower than predicted by WISC score
3. 34 *low achievers.* Performance was low, but not markedly lower than predicted from WISC

Testing Hypotheses about Reading Development

FSSSF(1996) gathered data annually on 363 six-year-olds until they reached the age of 16.

Besides completing the Woodcock-Johnson Psycho-educational Test Battery each year, they also completed the WISC every second year.

Three groups were identified:

1. 301 *normal readers*. Performance was within limits concomitant with WISC score
2. 28 *discrepant readers*. Performance was significantly lower than predicted by WISC score
3. 34 *low achievers*. Performance was low, but not markedly lower than predicted from WISC

Testing Hypotheses about Reading Development

Findings

- A quadratic model with a fixed asymptote modification was fit to each individual's data
- The two groups of disabled readers were statistically indistinguishable, but both differed significantly (about 30 points) from normal readers in their eventual plateau

Testing Hypotheses about Reading Development

Findings

- A quadratic model with a fixed asymptote modification was fit to each individual's data
- The two groups of disabled readers were statistically indistinguishable, but both differed significantly (about 30 points) from normal readers in their eventual plateau

Three Important Features of a Study of Change

Methodologically meaningful studies of change are characterized by 3 key features:

- Three or more waves of data
- An outcome whose values change systematically over time
- A sensible metric for measuring and recording time

Three Important Features of a Study of Change

Methodologically meaningful studies of change are characterized by 3 key features:

- Three or more waves of data
- An outcome whose values change systematically over time
- A sensible metric for measuring and recording time

Three Important Features of a Study of Change

Methodologically meaningful studies of change are characterized by 3 key features:

- Three or more waves of data
- An outcome whose values change systematically over time
- A sensible metric for measuring and recording time

Multiple Waves of Data

One or two waves of measurement are inadquate, because:

- The true shape of the change trajectory cannot be assessed with only two waves, because two points have only one interval, and so shape cannot be assessed. For example, consider the outcome values 1,8 assessed at times 1,2. These two data points perfectly fit the function $Y = X^3$ and the function $Y = 7X - 6$.
- Error variance cannot be distinguished from variance due to trajectory

In general, the more waves the better.

Multiple Waves of Data

One or two waves of measurement are inadquate, because:

- The true shape of the change trajectory cannot be assessed with only two waves, because two points have only one interval, and so shape cannot be assessed. For example, consider the outcome values 1,8 assessed at times 1,2. These two data points perfectly fit the function $Y = X^3$ and the function $Y = 7X - 6$.
- Error variance cannot be distinguished from variance due to trajectory

In general, the more waves the better.

A Sensible Time Metric

- Time is the fundamental independent variable in any study of change. It must be measured reliably, validly, and in an appropriate metric
- Coie, et al. used *grade* because they saw this as more meaningful in the context of antisocial behavior than chronological age
- Francis, et al. used *age*, because, among other reasons, the test scores were normed by age
- In many situations, you will have alternatives. For example, if you are studying longevity of automobiles, you'll consider age and mileage as possible time metrics

A Sensible Time Metric

- Time is the fundamental independent variable in any study of change. It must be measured reliably, validly, and in an appropriate metric
- Coie, et al. used *grade* because they saw this as more meaningful in the context of antisocial behavior than chronological age
- Francis, et al. used *age*, because, among other reasons, the test scores were normed by age
- In many situations, you will have alternatives. For example, if you are studying longevity of automobiles, you'll consider age and mileage as possible time metrics

A Sensible Time Metric

- Time is the fundamental independent variable in any study of change. It must be measured reliably, validly, and in an appropriate metric
- Coie, et al. used *grade* because they saw this as more meaningful in the context of antisocial behavior than chronological age
- Francis, et al. used *age*, because, among other reasons, the test scores were normed by age
- In many situations, you will have alternatives. For example, if you are studying longevity of automobiles, you'll consider age and mileage as possible time metrics

A Sensible Time Metric

- Time is the fundamental independent variable in any study of change. It must be measured reliably, validly, and in an appropriate metric
- Coie, et al. used *grade* because they saw this as more meaningful in the context of antisocial behavior than chronological age
- Francis, et al. used *age*, because, among other reasons, the test scores were normed by age
- In many situations, you will have alternatives. For example, if you are studying longevity of automobiles, you'll consider age and mileage as possible time metrics

A Sensible Time Metric

- Spacing can be equal or unequal
- Data can be *time-structured* (everyone is measured according to the same schedule) or *time-unstructured* (different individuals have different schedules).
- Data can be *balanced* (equal number of waves per individual) or *unbalanced* (different number of waves per individual)

A Sensible Time Metric

- Spacing can be equal or unequal
- Data can be *time-structured* (everyone is measured according to the same schedule) or *time-unstructured* (different individuals have different schedules).
- Data can be *balanced* (equal number of waves per individual) or *unbalanced* (different number of waves per individual)

A Sensible Time Metric

- Spacing can be equal or unequal
- Data can be *time-structured* (everyone is measured according to the same schedule) or *time-unstructured* (different individuals have different schedules).
- Data can be *balanced* (equal number of waves per individual) or *unbalanced* (different number of waves per individual)

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
    - Ideally the measure should be the same
    - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
    - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
  - Ideally the measure should be the same
  - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
  - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
    - Ideally the measure should be the same
    - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
    - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
  - Ideally the measure should be the same
  - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
  - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
  - Ideally the measure should be the same
  - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
  - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
  - Ideally the measure should be the same
  - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
  - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Continuous, Time-Varying, Interval Scale Outcome

- The substantive construct chosen as the outcome depends, of course, on the research question of interest
- Psychometric qualities of reliability and validity are paramount
- A good case must be sustainable for *time-equivalence* of the measure. That is, equal scores should have an equal meaning at different times
  - Ideally the measure should be the same
  - Simply standardizing won't necessarily help
- A measure may be valid at Time 1, but not valid at Time 2.
  - For example, (Lord, 1963) multiplication skill might be a valid measure of mathematical ability for 3rd graders, but more a measure of memory for 8th graders

Exploring Longitudinal Data

In this section, we examine

1. How to set up longitudinal data for optimal analysis
2. How to explore longitudinal data graphically, and with summary statistics

Much of the example code that follows was presented on the UCLA Statistics page dedicated to the text by Singer and Willett (2003).

Exploring Longitudinal Data

In this section, we examine

1. How to set up longitudinal data for optimal analysis
2. How to explore longitudinal data graphically, and with summary statistics

Much of the example code that follows was presented on the UCLA Statistics page dedicated to the text by Singer and Willett (2003).

Organizing Longitudinal Data – The NYS Example

Singer and Willett (2003), in their outstanding text *Applied Longitudinal Data Analysis*, discuss an example from the National Youth Survey (NYS) in their introductory treatment.

- There were 5 waves of data, recorded when participants were 11,12,13,14, and 15 years of age
- Each year, participants filled out a 9-item instrument designed to assess their tolerance of deviant behavior
- The items were on a 4-points scale, with 1 = very wrong, 2 = wrong, 3 = a little bit wrong, 4 = not wrong at all

Organizing Longitudinal Data – The NYS Example

Singer and Willett (2003), in their outstanding text *Applied Longitudinal Data Analysis*, discuss an example from the National Youth Survey (NYS) in their introductory treatment.

- There were 5 waves of data, recorded when participants were 11,12,13,14, and 15 years of age
- Each year, participants filled out a 9-item instrument designed to assess their tolerance of deviant behavior
- The items were on a 4-points scale, with 1 = very wrong, 2 = wrong, 3 = a little bit wrong, 4 = not wrong at all

Organizing Longitudinal Data – The NYS Example

Singer and Willett (2003), in their outstanding text *Applied Longitudinal Data Analysis*, discuss an example from the National Youth Survey (NYS) in their introductory treatment.

- There were 5 waves of data, recorded when participants were 11,12,13,14, and 15 years of age
- Each year, participants filled out a 9-item instrument designed to assess their tolerance of deviant behavior
- The items were on a 4-points scale, with 1 = very wrong, 2 = wrong, 3 = a little bit wrong, 4 = not wrong at all

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than \$5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than \$50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than \$5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than \$50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

The behaviors surveyed by the NYS were:

- Cheat on tests
- Purposely destroy the property of others
- Use marijuana
- Steal something worth less than $5
- Hit or threaten someone without reason
- Use alcohol
- Break into a building or vehicle to steal
- Sell hard drugs
- Steal something worth more than $50

On each occasion, the tolerance score (TOL) was simply the average of the responses on the 9 items.

The NYS Survey

In addition, potential predictors were recorded. These predictors are classified as *time-invariant* since their values do not change across the time period of the study. The predictors are MALE and EXPOSURE. The latter was a measure of prior exposure to deviant behaviors, and was taken at the age of 11.

Person-Level Data

Singer and Willett (2003) discuss two fundamental forms of data file. "Person-Level" data includes one row of data per person, with all variables (including measurement at all waves) recorded in that row.

Let's load in the file and take a look.

```
> tolerance <- read.table("tolerance1.txt",header=T,sep=",")
> print(tolerance)
```

```
     id tol11 tol12 tol13 tol14 tol15 male exposure
1     9  2.23  1.79  1.90  2.12  2.66    0     1.54
2    45  1.12  1.45  1.45  1.45  1.99    1     1.16
3   268  1.45  1.34  1.99  1.79  1.34    1     0.90
4   314  1.22  1.22  1.55  1.12  1.12    0     0.81
5   442  1.45  1.99  1.45  1.67  1.90    0     1.13
6   514  1.34  1.67  2.23  2.12  2.44    1     0.90
7   569  1.79  1.90  1.90  1.99  1.99    0     1.99
8   624  1.12  1.12  1.22  1.12  1.22    1     0.98
9   723  1.22  1.34  1.12  1.00  1.12    0     0.81
10  918  1.00  1.00  1.22  1.99  1.22    0     1.21
11  949  1.99  1.55  1.12  1.45  1.55    1     0.93
12  978  1.22  1.34  2.12  3.46  3.32    1     1.59
13 1105  1.34  1.90  1.99  1.90  2.12    1     1.38
14 1542  1.22  1.22  1.99  1.79  2.12    0     1.44
15 1552  1.00  1.12  2.23  1.55  1.55    0     1.04
16 1653  1.11  1.11  1.34  1.55  2.12    0     1.25
```

Person-Period Data

"Person-Period" data contains one row for each period a person is measured in.

Let's load in the file and see what I mean.

```
> tolerance.pp <- read.table("tolerance1_pp.txt", sep=",", header=T)
> print(tolerance.pp)
```

```
    id age tolerance male exposure time
1    9  11      2.23    0     1.54    0
2    9  12      1.79    0     1.54    1
3    9  13      1.90    0     1.54    2
4    9  14      2.12    0     1.54    3
5    9  15      2.66    0     1.54    4
6   45  11      1.12    1     1.16    0
7   45  12      1.45    1     1.16    1
8   45  13      1.45    1     1.16    2
9   45  14      1.45    1     1.16    3
10  45  15      1.99    1     1.16    4
11 268  11      1.45    1     0.90    0
12 268  12      1.34    1     0.90    1
13 268  13      1.99    1     0.90    2
14 268  14      1.79    1     0.90    3
15 268  15      1.34    1     0.90    4
```

Disadvantages of the Person-Level Format

Although the Person-Level format makes it easy to plot an individual's empirical growth record, it has 4 distinct disadvantages

- It leads naturally to noninformative summaries
  - For example, although it is natural to compute the correlation matrix between the tolerance variables, this may not be very informative
- It omits an explicit time variable
- It can be very inefficient when the number and/or spacing of waves vary across individuals
- It cannot easily handle time-varying predictors

Disadvantages of the Person-Level Format

Although the Person-Level format makes it easy to plot an individual's empirical growth record, it has 4 distinct disadvantages

- It leads naturally to noninformative summaries
  - For example, although it is natural to compute the correlation matrix between the tolerance variables, this may not be very informative
- It omits an explicit time variable
- It can be very inefficient when the number and/or spacing of waves vary across individuals
- It cannot easily handle time-varying predictors

Disadvantages of the Person-Level Format

Although the Person-Level format makes it easy to plot an individual's empirical growth record, it has 4 distinct disadvantages

- It leads naturally to noninformative summaries
  - For example, although it is natural to compute the correlation matrix between the tolerance variables, this may not be very informative
- It omits an explicit time variable
- It can be very inefficient when the number and/or spacing of waves vary across individuals
- It cannot easily handle time-varying predictors

Disadvantages of the Person-Level Format

Although the Person-Level format makes it easy to plot an individual's empirical growth record, it has 4 distinct disadvantages

- It leads naturally to noninformative summaries
    - For example, although it is natural to compute the correlation matrix between the tolerance variables, this may not be very informative
- It omits an explicit time variable
- It can be very inefficient when the number and/or spacing of waves vary across individuals
- It cannot easily handle time-varying predictors

Disadvantages of the Person-Level Format

Although the Person-Level format makes it easy to plot an individual's empirical growth record, it has 4 distinct disadvantages

- It leads naturally to noninformative summaries
  - For example, although it is natural to compute the correlation matrix between the tolerance variables, this may not be very informative
- It omits an explicit time variable
- It can be very inefficient when the number and/or spacing of waves vary across individuals
- It cannot easily handle time-varying predictors

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Advantages of the Person-Period Format

The Person-Period format is characterized by

1. An identifier
2. A time indicator
3. Outcome variable(s)
4. Predictor variable(s)

It offers a number of advantages:

1. Easy recording of outcome variables
2. Easy recording of time-varying predictors
3. Explicit time variable

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
   1. Least squares regression line, or
   2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
   1. Residual variance of individual plots
   2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
   1. Least squares regression line, or
   2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
   1. Residual variance of individual plots
   2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
   1. Least squares regression line, or
   2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
   1. Residual variance of individual plots
   2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
    1. Least squares regression line, or
    2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
    1. Residual variance of individual plots
    2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
    1. Least squares regression line, or
    2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
    1. Residual variance of individual plots
    2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
    1. Least squares regression line, or
    2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
    1. Residual variance of individual plots
    2. Regression parameter estimates (slopes and intercepts)

Descriptive Analysis of Individual Change

Singer and Willett discuss several basic exploratory techniques for getting the "feel" of your data, including:

1. Empirical Growth Plots, augmented by
   1. Least squares regression line, or
   2. Nonparametric smoothed fit
2. R^2 statistics for individual growth plots
3. Stem leaf diagrams of:
   1. Residual variance of individual plots
   2. Regression parameter estimates (slopes and intercepts)

Trellis Plot of Individual Growth Curves

```
> # load lattice library for xyplot function
> library(lattice)
> xyplot(tolerance ~ age | id,
+ data=tolerance.pp, as.table=T)
> update(trellis.last.object(),
+  strip = strip.custom(strip.names = TRUE,
+  strip.levels = TRUE))
```

Trellis Plot of Individual Growth Curves

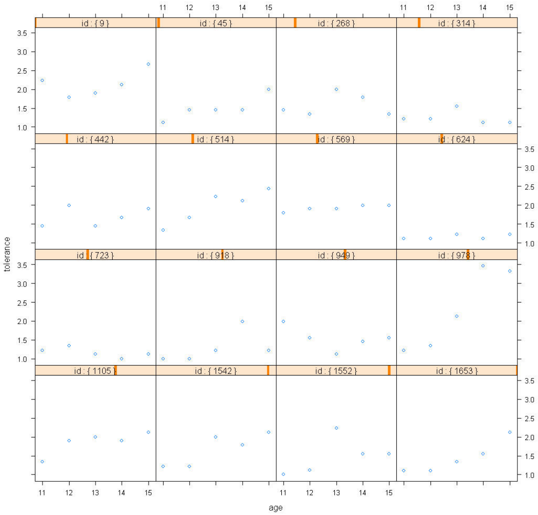

Adding Nonparametric Smoothing

Often it can enhance readability to add a fit line to the individual growth plots. Here we add a loess smooth line.

```
> xyplot(tolerance~age | id, data=tolerance.pp,
+       prepanel = function(x, y)
+       prepanel.loess(x, y, family="gaussian"),
+       xlab = "Age", ylab = "Tolerance",
+       panel = function(x, y) {
+       panel.xyplot(x, y)
+       panel.loess(x,y, family="gaussian") },
+       ylim=c(0, 4), as.table=T)
> update(trellis.last.object(),
+  strip = strip.custom(strip.names = TRUE,
+  strip.levels = TRUE))
```

Adding Nonparametric Smoothing

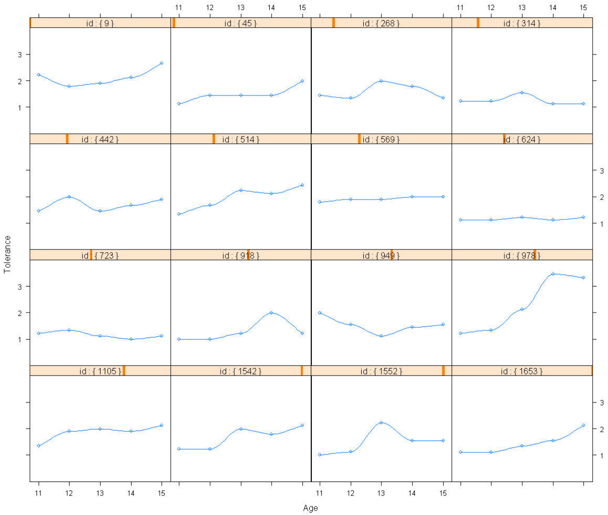

Computing Linear Fit Lines

You can compute linear fit lines for all the individual growth curves in one R command.

The output is extensive so we exclude it here.

```
> attach(tolerance.pp)
> by(tolerance.pp, id,
+ function(x) summary(lm(tolerance ~ time, data=x)))
```

Stem-Leaf Plot for the Intercepts

```
The following object(s) are masked _by_ .GlobalEnv :

 tolerance

> int <- by(tolerance.pp, id,
+ function(data) coefficients(lm(tolerance ~ time, data = data))[[1]])
> int <- unlist(int)
> names(int) <- NULL
> summary(int)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 0.954   1.140   1.290   1.360   1.550   1.900

> stem(int, scale=2)

 The decimal point is 1 digit(s) to the left of the |

  9 | 5
 10 | 03
 11 | 2489
 12 | 7
 13 | 1
 14 | 3
 15 | 448
 16 |
 17 | 3
 18 | 2
 19 | 0
```

Stem-Leaf Plot for the Slopes

```
> rate <- by(tolerance.pp, id,
+ function(data) coefficients(lm(tolerance ~ time, data = data))[[2]])
> rate <- unlist(rate)
> names(rate) <- NULL
> summary(rate)

   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-0.0980  0.0223  0.1310  0.1310  0.1900  0.6320

> stem(rate, scale=2)

  The decimal point is 1 digit(s) to the left of the |

  -1 | 0
  -0 | 53
   0 | 2256
   1 | 24567
   2 | 457
   3 |
   4 |
   5 |
   6 | 3
```

Stem-Leaf Plot for the R^2 Values

```
> rsq <- by(tolerance.pp, id,
+ function(data)summary(lm(tolerance ~ time, data = data))$r.squared)
> rsq <- unlist(rsq)
> names(rsq) <- NULL
> summary(rsq)

   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 0.0154  0.2500  0.3920  0.4910  0.7970  0.8860

> stem(rsq, scale=2)

  The decimal point is 1 digit(s) to the left of the |

  0 | 27
  1 | 3
  2 | 55
  3 | 113
  4 | 5
  5 |
  6 | 8
  7 | 78
  8 | 6889
```

Trellis Plot of Linear Fit Lines

```
> attach(tolerance)
> xyplot(tolerance ~ age | id, data=tolerance.pp,
+   panel = function(x, y){
+     panel.xyplot(x, y)
+     panel.lmline(x, y)
+   }, ylim=c(0, 4), as.table=T)
> update(trellis.last.object(),
+  strip = strip.custom(strip.names = TRUE,
+  strip.levels = TRUE))
```

Trellis Plot of Linear Fit Lines

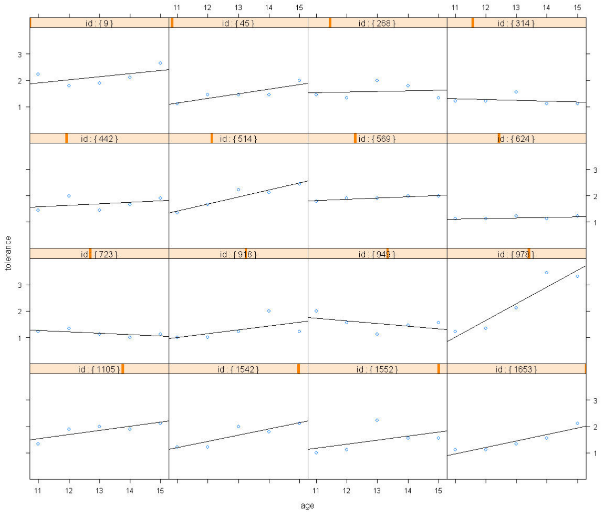

Examining the Entire Set of Smooth Trajectories

One effective way of exploring interindividual differences in change is to plot, on a single graph, the entire set of fits for the individual trajectories. Here is a plot of the raw data.

```
> interaction.plot(tolerance.pp$age,
+ tolerance.pp$id, tolerance.pp$tolerance)
```

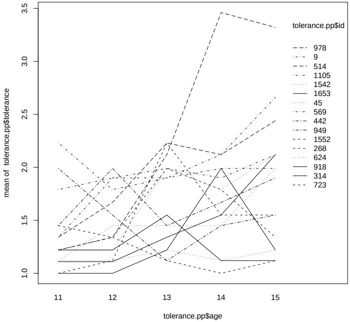

Examining the Entire Set of Smooth Trajectories

A cleaner picture is presented by plotting together the OLS linear fit lines.

```
> # fitting the linear model by id
> fit <- by(tolerance.pp, id,
+  function(bydata) fitted.values(lm(tolerance ~ time, data=bydata)))
> fit <- unlist(fit)
> # plotting the linear fit by id
> interaction.plot(age, id, fit, xlab="age", ylab="tolerance")
```

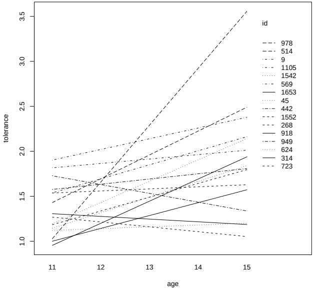

Summary Statistics for the Set of Trajectory Coefficients

Computing summary statistics for the set of trajectory coefficients can provide useful information about the overall trends in the data.

Summary Statistics for the Set of Trajectory Coefficients

```
> #obtaining the intercepts from linear model by id
> ints <- by(tolerance.pp, tolerance.pp$id,
+ function(data) coefficients(lm(tolerance ~ time, data=data))[[1]])
> ints1 <- unlist(ints)
> names(ints1) <- NULL
> mean(ints1)

[1] 1.358

> sqrt(var(ints1))

[1] 0.2978

> #obtaining the slopes from linear model by id
> slopes <- by(tolerance.pp, tolerance.pp$id,
+  function(data) coefficients(lm(tolerance ~ time, data=data))[[2]])
> slopes1 <- unlist(slopes)
> names(slopes1) <- NULL
> mean(slopes1)

[1] 0.1308

> sqrt(var(slopes1))

[1] 0.1723

> cor( ints1, slopes1)

[1] -0.4481
```

Plotting Trajectories by Level of Potential Predictor

In assessing whether a potential predictor actually predicts differential change, a good place to start is by plotting individual growth trajectories as a function of level of a potential predictor.

Male Trajectories

We begin by plotting trajectories separately for males and females. First the male plot.

```
> # fitting the linear model by id, males only
> tolm <- tolerance.pp[male==0 , ]
> fitmlist <- by(tolm, tolm$id,
+ function(bydata) fitted.values(lm(tolerance ~ time, data=by
> fitm <- unlist(fitmlist)
> #appending the average for the whole group
> lm.m <- fitted( lm(tolerance ~ time, data=tolm) )
> names(lm.m) <- NULL
> fit.m2 <- c(fitm, lm.m[1:5])
> age.m <- c(tolm$age, seq(11,15))
> id.m <- c(tolm$id, rep(111, 5))
```

Male Trajectories

```
> interaction.plot(age.m, id.m, fit.m2, ylim=c(0, 4), xlab="AGE", ylab="TOLERANCE", lwd=1)
> title(main="Males")
```

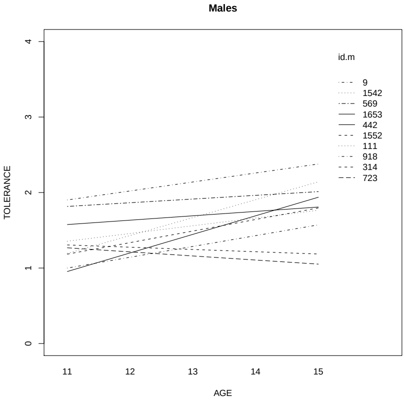

Female Trajectories

We then compare them with the female trajectories.

```
> #fitting the linear model by id, females only
> tol.pp.fm <- tolerance.pp[tolerance.pp$male==1 , ]
> fit.fm <- by(tol.pp.fm, tol.pp.fm$id,
+  function(data) fitted.values(lm(tolerance ~ time,
+  data=data)))
> fit.fm1 <- unlist(fit.fm)
> names(fit.fm1) <- NULL
> #appending the average for the whole group
> lm.fm <- fitted( lm(tolerance ~ time, data=tol.pp.fm) )
> names(lm.fm) <- NULL
> fit.fm2 <- c(fit.fm1, lm.fm[1:5])
> age.fm <- c(tol.pp.fm$age, seq(11,15))
> id.fm <- c(tol.pp.fm$id, rep(1111, 5))
```

Female Trajectories

```
> interaction.plot(age.fm, id.fm, fit.fm2, ylim=c(0, 4), xlab="AGE", ylab="TOLERANCE", lwd=
> title(main="Females")
```

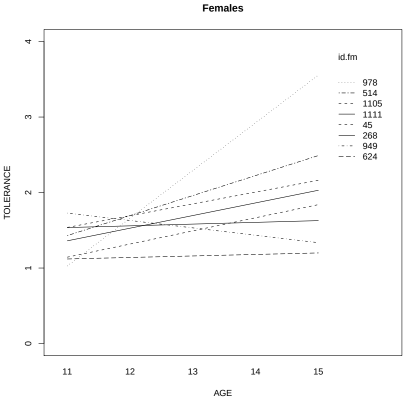

Substantive Conclusion

Certainly, female ID 978 appears to be an unusual case.
Overall, there appears to be little difference between males and females in terms of the overall pattern of the trajectories.

Low Exposure Trajectories

Next we compare low exposure subjects to those with high exposure. First the low exposure plot.

```
> #fitting the linear model by id, low exposure
> tol.pp.low <- tolerance.pp[tolerance.pp$exposure < 1.145 , ]
> fit.low <- by(tol.pp.low, tol.pp.low$id,
+  function(data) fitted.values(lm(tolerance ~ time,
+  data=data)))
> fit.low1 <- unlist(fit.low)
> names(fit.low1) <- NULL
> #appending the average for the whole group
> lm.low <- fitted( lm(tolerance ~ time, data=tol.pp.low) )
> names(lm.low) <- NULL
> fit.low2 <- c(fit.low1, lm.low[1:5])
> age.low <- c(tol.pp.low$age, seq(11,15))
> id.low <- c(tol.pp.low$id, rep(1, 5))
```

Low Exposure Trajectories

```
> interaction.plot(age.low, id.low, fit.low2, ylim=c(0, 4),
+ xlab="AGE", ylab="TOLERANCE", lwd=1)
> title(main="Low Exposure")
```

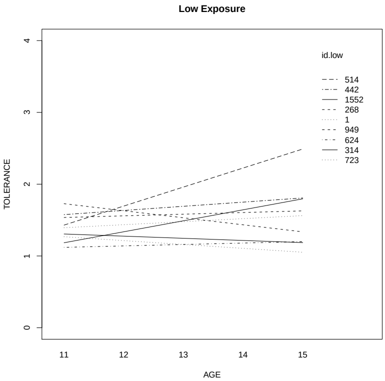

High Exposure Trajectories

Next the plot for high exposure subjects.

```
> #fitting the linear model by id, high exposure
> tol.pp.hi <- tolerance.pp[tolerance.pp$exposure >= 1.145 , ]
> fit.hi <- by(tol.pp.hi, tol.pp.hi$id,
+  function(data) fitted.values(lm(tolerance ~ time,
+  data=data)))
> fit.hi1 <- unlist(fit.hi)
> names(fit.hi1) <- NULL
> #appending the average for the whole group
> lm.hi <- fitted( lm(tolerance ~ time, data=tol.pp.hi) )
> names(lm.hi) <- NULL
> fit.hi2 <- c(fit.hi1, lm.hi[1:5])
> age.hi <- c(tol.pp.hi$age, seq(11,15))
> id.hi <- c(tol.pp.hi$id, rep(1, 5))
```

High Exposure Trajectories

```
> interaction.plot(age.hi, id.hi, fit.hi2, ylim=c(0, 4),
+ xlab="AGE", ylab="TOLERANCE", lwd=1)
> title(main="High Exposure")
```

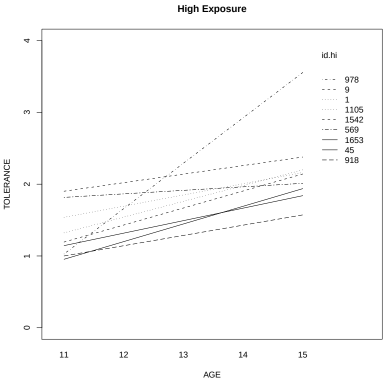

Substantive Conclusion

Even discounting ID 978, there appears to be an overall trajectory difference between individuals with low exposure and those with high exposure. Although there is not much difference in initial level (intercept), there appears to be a substantial difference in slope.

Therefore, it seems those with high exposure gain tolerance more rapidly as they age.

Plotting Fitted Intercepts by Sex

```
> plot(tolerance$male, ints1, xlab="Male",
+ ylab="Fitted initial status",
+       xlim=c(0, 1), ylim=c(0.5, 2.5))
> cor(tolerance$male, ints1)

[1] 0.00863
```

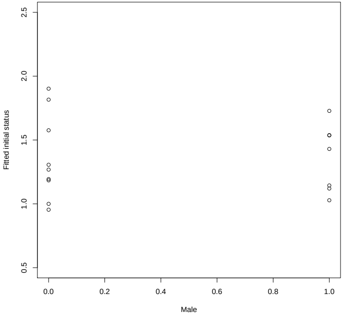

Plotting Fitted Slopes by Sex

```
> plot(tolerance$male, slopes1, xlab="Male",
+  ylab="Fitted rate of change",
+  xlim=c(0, 1), ylim=c(-0.4, .8))
> cor(tolerance$male, slopes1)

[1] 0.1936
```

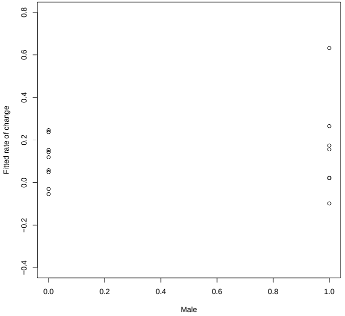

Plotting Fitted Intercepts by Exposure

```
> plot(tolerance$exposure, ints1,
+ xlab="Exposure", ylab="Fitted initial status",
+  xlim=c(0, 2.5), ylim=c(0.5, 2.5))
> cor(tolerance$exposure, ints1)

[1] 0.2324
```

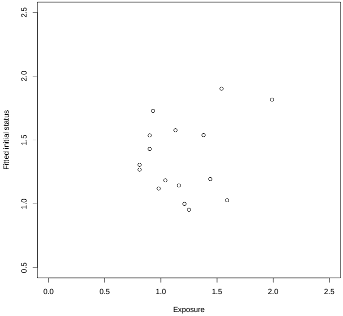

Plotting Fitted Slopes by Exposure

```
> plot(tolerance$exposure, slopes1,
+ xlab = "Exposure", ylab =
+ "Fitted rate of change",
+ xlim = c(0, 2.5), ylim = c(-0.2, 0.8))
> cor(tolerance$exposure, slopes1)

[1] 0.4421
```