Analysis of Interrupted Time Series with Segmented Regression

Nancy A. Perrin, PhD
Center for Health Research

October 2, 2009

Center for Health Research

Overview

- The basic time series design
- Why it is a strong design
  - Review of quasi-experimental designs
  - Rubin's Causal Model
- Segmented Regression
  - Data structure
  - Models
  - Interpretation
- Enhancing the interrupted time series design

Center for Health Research

The Situation

- Data collected at **equally spaced intervals** on some variable over time
  - Number of kids caught smoking at a high school
    - Monthly proportion of students caught smoking
- Start a new program, intervention, or some change that effects the variable
  - Campaign to reduce smoking aimed at students, parents, and teachers
- Continue to collect the variable at the same equally spaced intervals over time
  - Pre and post implementation of the intervention

Center for Health Research

The Approach

- Estimate the trend in the dependent variable prior to the intervention
- Estimate the trend after the intervention
- Test for changes in dependent variable pre and post the intervention
- Test for changes in the slope of the trend pre and post intervention

Center for Health Research

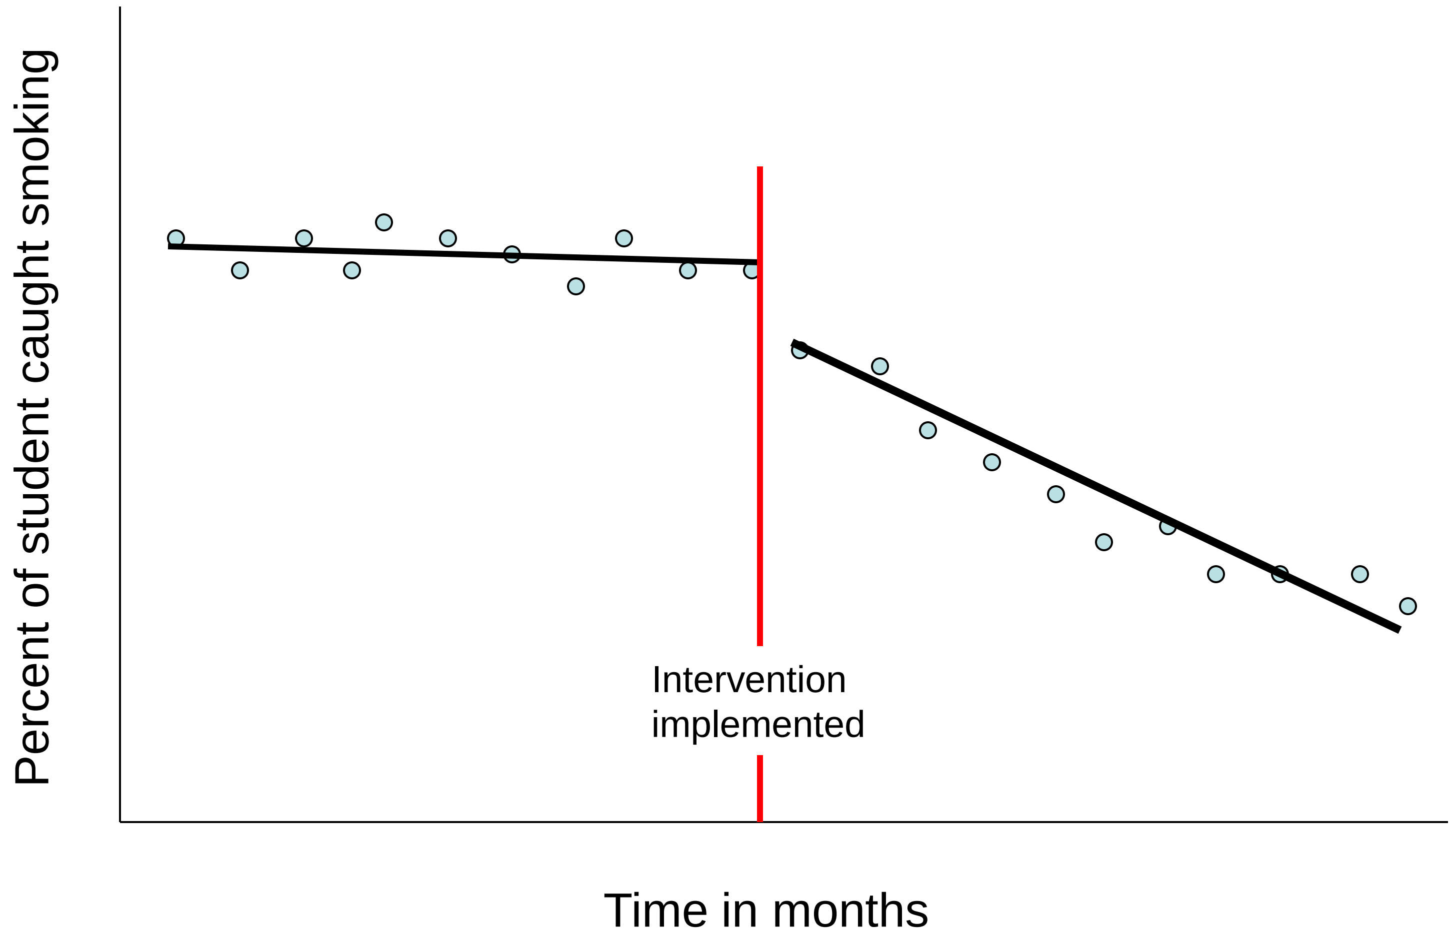

Center for Health Research

Strengths of the time series design

- Excellent method of conducting naturalistic studies of the effects of a change at a systems level*
- Intuitive graphical representation
- Takes the trend before the intervention into account
- Can have good statistical power with relatively few data points if the trend prior to the intervention is fairly stable
- Can capitalize on existing data
- Size of the effect can be estimated at different times post the intervention

*Wes, Biesanz, Pitts (2000)

Center for Health Research

But wait ….

- There is only one group being observed
  - No control group
- There is no random assignment
- What if the change “just happened” at the same time as the intervention
- This really seems like a weak design
- Let’s take a closer look
  - But first we need a little review on quasi-experimental designs

Center for Health Research

Strong Alternatives to the Randomized Experiment*

- Regression discontinuity design
  - Assignment to naturally occurring groups based on a continuous variable
    - Assigned based on whether or not the subject exceeds some threshold
- Interrupted time series
  - Outcome is collected at equally spaced intervals over a long period of time
  - Intervention/event occurs at a specific time point
    - Test for changes in the level and slope of the time series pre and post intervention

*Shadish, Cook, Campbell (2002); Reichardt & Mark (1997); Marcantonio & Cook (1994)

Center for Health Research

Rubin's Causal Model*

- Helpful in identifying strengths and limitations of different designs
- Causal Effect
  - Causality is at the level of a single person
    - X caused Y in a given case
  - Difference between what would have happened to the subject under the treatment condition and what would have happened to the same subject under the control condition
    - Can never observe this, only infer it

*Rubin 1974, 1978, 1986

Randomization

- Used to equate groups at pre-test
- Group assignment can be expected on average to be independent of pre-test scores
  - Groups have equivalent means on measured & unmeasured variables at pre-test
  - Group assignment in uncorrelated with any pre-test variables
- Randomization replaces definitive statements about causal effects with probabilistic statements
  - No longer at the level of a subject
  - "On average" the two groups will not differ at pre-test
  - If X occurs the probability that Y occurs is…

Unit of Homogeneity

- If the groups are assumed to be equal, then it does not matter which group received treatment
  - Basic idea behind randomization
  - Place we want to get to in quasi-experimental designs

Center for Health Research

Regression Discontinuity Design

- Treatment/control is assigned based on some continuous measure
  - Those above a certain score get treatment
    - BP above some level → go on medication
- Outcome measured after some amount of time
- Treatment effect estimated by

$$Y = b_o + b_1X + b_2D$$

X is score on continuous measure used to determine who got treatment

D is dummy variable for group assignment

Center for Health Research

Hmmmm.....

- The groups are not equivalent at pre-test
- The group effect is estimated conditioned on the variable used to assign group membership
  - If this variable is the sole basis on which participants are assigned to groups then we have an unbiased estimate for the group effect*
    - Unbiased estimate of the difference between the mean of the two groups on the outcome
      - Adjusted mean difference

*Rubin 1977

Center for Health Research

Interrupted Time Series Design

- Observe some measure at regular intervals over time
- Note there is one observation per time point
- At some point an interruption occurs
  - New program, change in guidelines, intervention

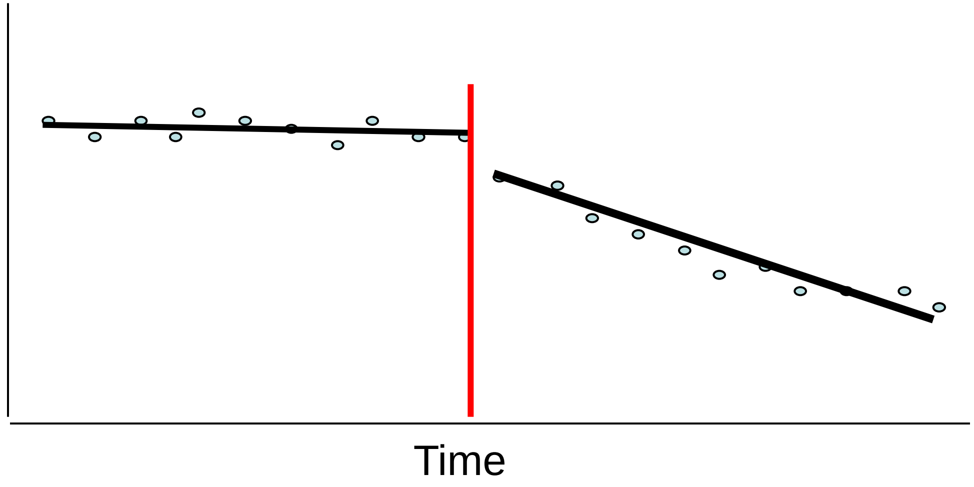

Center for Health Research

Rubin Causal Model View of Interrupted Time Series

- Expected value of the treatment group (post-interruption) is compared to the expected value of the control group (pre-interruption)
  - conditioned on the specific point in time the change was introduced
- **If** …
  - time is a good proxy for the rule in which the intervention and treatment groups were assigned
  - the correct functional form of the relationship between time and the outcome is specified
    - Linear, quadratic

Center for Health Research

Then…

- Controlling for time will lead to an unbiased estimate of the treatment effect

$$Y = b_o + b_1X + b_2D$$

X score on continuous measure used to determine who got treatment

→ in this case it is time

D dummy variable for group assignment

→ in this case it is before or after the intervention was implemented

Center for Health Research

Why in times series a strong alternative to randomization?

- Can get an unbiased estimate of the difference between pre and post intervention
  - Because time is the rule on which it was decided who got the intervention
  - Time must be included in the model

Analysis of Time Series Data: Segmented Regression

- Basic model

$$Y_t = b_o + b_1 T + b_2 D + b_3 P + e_t$$

- T is time from the start of the observational period
  - Continuous variable beginning at 1
- D is a dummy variable for pre or post intervention
  - Coded 0 prior to intervention, 1 post intervention
- P is time since the intervention
  - Prior to intervention - coded 0
  - Post intervention – continuous starting at 1
- e_t is the random variation at time *t* not explained by the model

Center for Health Research

Note: There is a single observation per time point

Obs	%smoking	T	D	P
1	62.5	1	0	0
2	65.8	2	0	0
3	61.4	3	0	0
4	62.3	4	0	0
5	61.8	5	0	0
6	64.6	6	0	0
7	63.1	7	0	0
8	63.8	8	0	0
9	55.4	9	1	1
10	52.1	10	1	2
11	51.7	11	1	3
12	53.4	12	1	4
13	50.1	13	1	5
14	48.3	14	1	6
15	45.9	15	1	7
16	40.2	16	1	8
17	43.8	17	1	9
18	38.7	18	1	10
19	37.5	19	1	11

Analysis

- Regression
  - Dependent variable – outcome of interest
    - Scores on the monthly observations
  - 3 independent variables
    - T – time
    - D – dummy for pre-post intervention
    - P – time since the intervention
- Ordinary Least Squares if model assumptions are met
- Autoregressive or ARIMA if you have serial dependency

Center for Health Research

$Y_t = b_o + b_1T + b_2D + b_3P$

- b_o is the baseline level of the outcome
  - Value at time zero
- b_1 is the slope prior to the intervention
  - Change over time before the intervention was implemented
- b_2 is the change in level immediately after the intervention
  - Change in the outcome measure from the last time point before the intervention to the first time point after the intervention
- b_3 is the change in the slope from pre to post intervention
  - Difference in the slope of the time period before the intervention and the slope of the time period after the intervention

Center for Health Research

$Y_t = b_o + b_1 T + b_2 D + b_3 P$

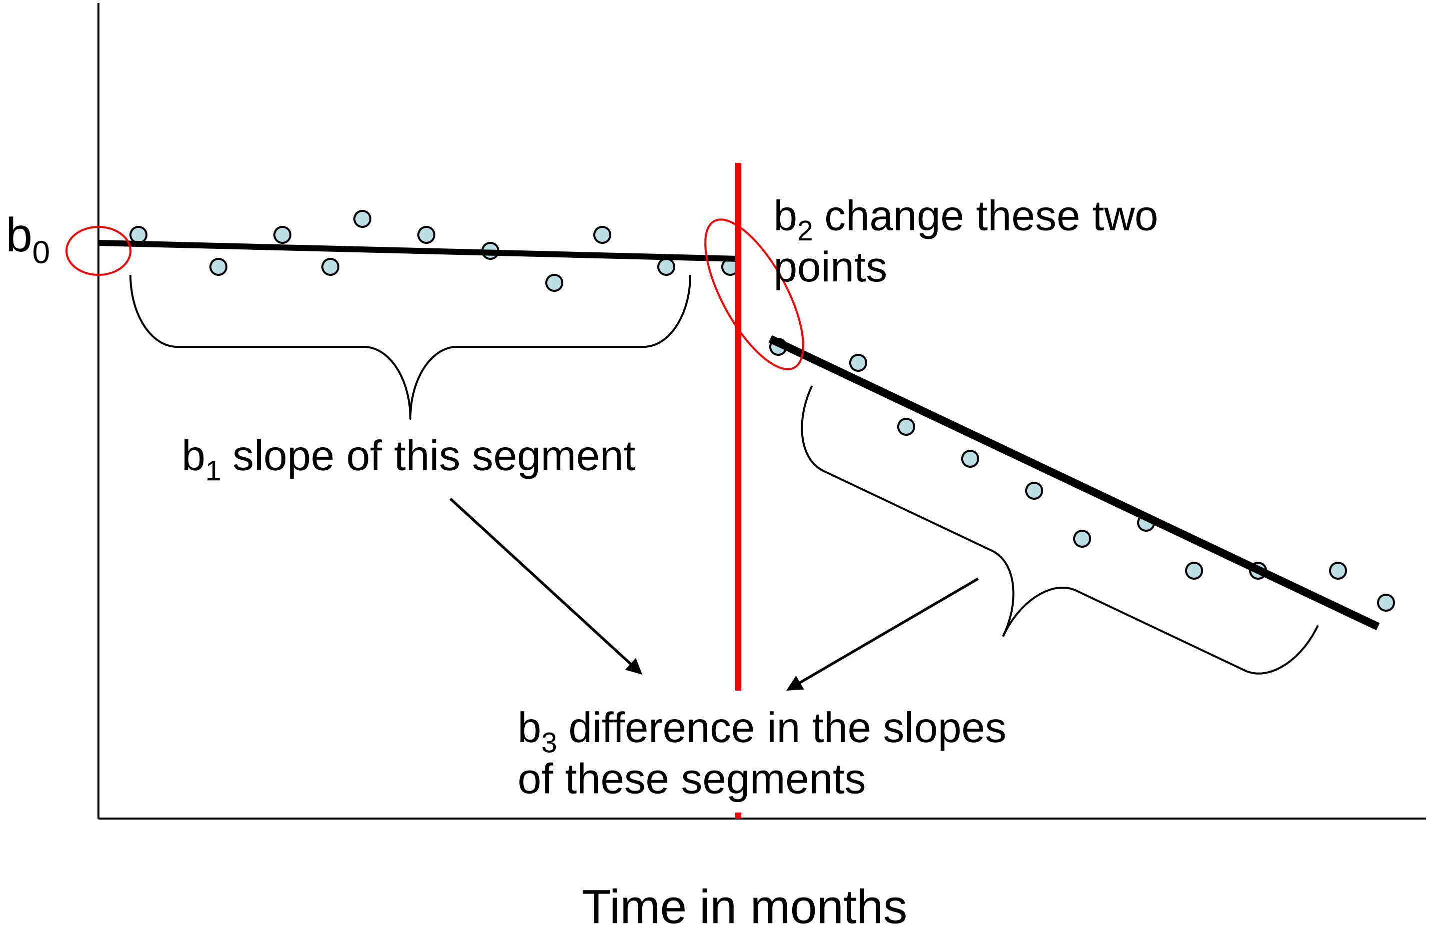

Center for Health Research

$Y_t = b_o + b_1T + b_2D + b_3P$

- When testing the change in level and slope, we are controlling for baseline trend
  - T is the trend prior to the intervention
  - It is also the rule that determines who got the intervention
- Testing changes in trend and level taking the baseline pattern into account
  - Difference over time are conditioned on the rule that determined who got the intervention

Center for Health Research

What is $b_1 + b_3$?

- b_1 is the slope prior to the intervention
- b_3 is the change in the slope from pre to post intervention

Center for Health Research

Fake Example: Smoking

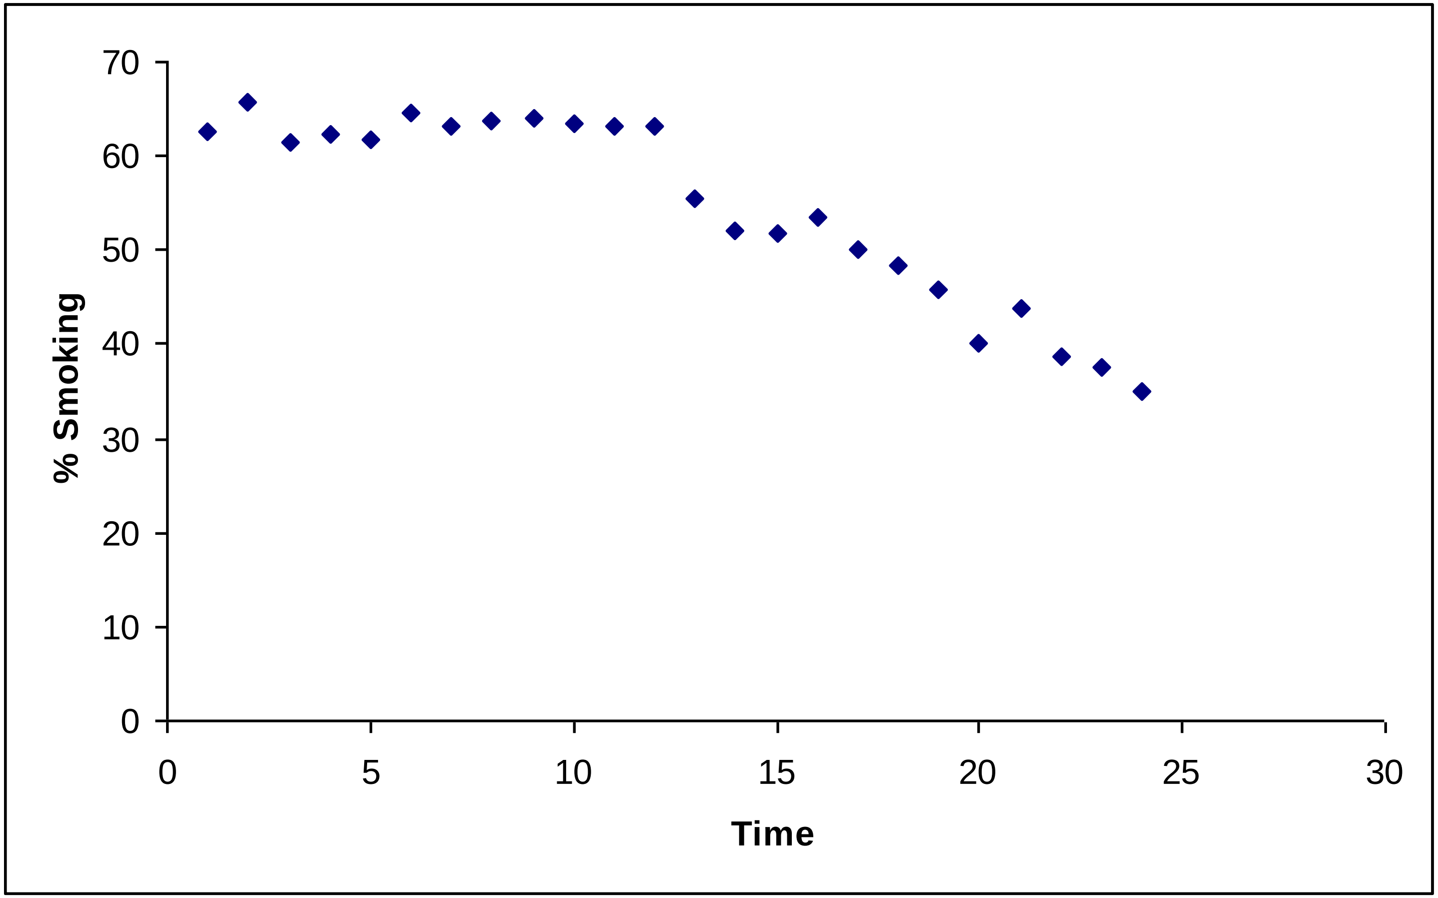

Center for Health Research

SPSS Output: OLS Regression

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	62.986	.982		64.129	.000
	T	.044	.133	.031	.333	.743
	D	-5.519	1.310	-.280	-4.214	.000
	P	-1.891	.189	-.780	-10.017	.000

a. Dependent Variable: % smoking

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.989[a]	.978	.975	1.5959

a. Predictors: (Constant), P, D, T

Center for Health Research

Interpretation

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	62.986	.982		64.129	.000
	T	.044	.133	.031	.333	.743
	D	-5.519	1.310	-.280	-4.214	.000
	P	-1.891	.189	-.780	-10.017	.000

a. Dependent Variable: % smoking

Trend prior to intervention was flat, not changing

There was a significant drop in % smoking after the intervention

There was a change in trend after the intervention, with a stronger decline over time in % smoking after the intervention

Center for Health Research

Interpreation

- What was the change in % smoking immediately after the campaign started?
- →It decreased by 5.5%
- What was the rate of change in smoking prior to the campaign?
- → Increasing by 0.04% per month
- What was the rate of change after the campaign started?
- → Decreasing by 1.85% per month
  - 0.044-1.891=1.847

Center for Health Research

Alternative Ways of Expressing Effects

- Can compare post intervention values estimated from the level and trend of the pre-period with observed post-intervention values
  - What would the trend look like in the absence of the intervention?
  - Counterfactual value
- Effect can be expressed as
  - Absolute difference between predicted outcome based on the intervention and counterfactual value
  - Ratio of the predicted and counterfactual values
- Need to select a point in time
  - 6 months post intervention, 12 months post intervention

Estimating Effects

$$Y_t = b_o + b_1 T + b_2 D + b_3 P$$

Value of outcome at 6 months post intervention

$$Y_t = b_o + b_{1*}18 + b_2*1 + b_3*6$$

Value of outcome at 6 months post intervention if the intervention had not happened

$$Y_t = b_o + b_{1*}18 + b_2*0 + b_3*0$$

Had a 12 month pre period

Calculating absolute difference and relative change

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	62.986	.982		64.129	.000
	T	.044	.133	.031	.333	.743
	D	-5.519	1.310	-.280	-4.214	.000
	P	-1.891	.189	-.780	-10.017	.000

a. Dependent Variable: % smoking

Y = 62.99 + (.04*18) – (5.52*1) – (1.89*6) = 46.85% smoking with intervention

Y = 62.99 + (.04*18) – (5.52*0) – (1.89*0) = 63.71% predicted to be smoking if the intervention had not happened

16.86% (63.71-46.85) absolute difference

26.46% (16.86/63.71*100) relative percent change

Can add confidence intervals around these numbers using the confidence intervals around the regression coefficient estimates

Center for Health Research

Smoking Example

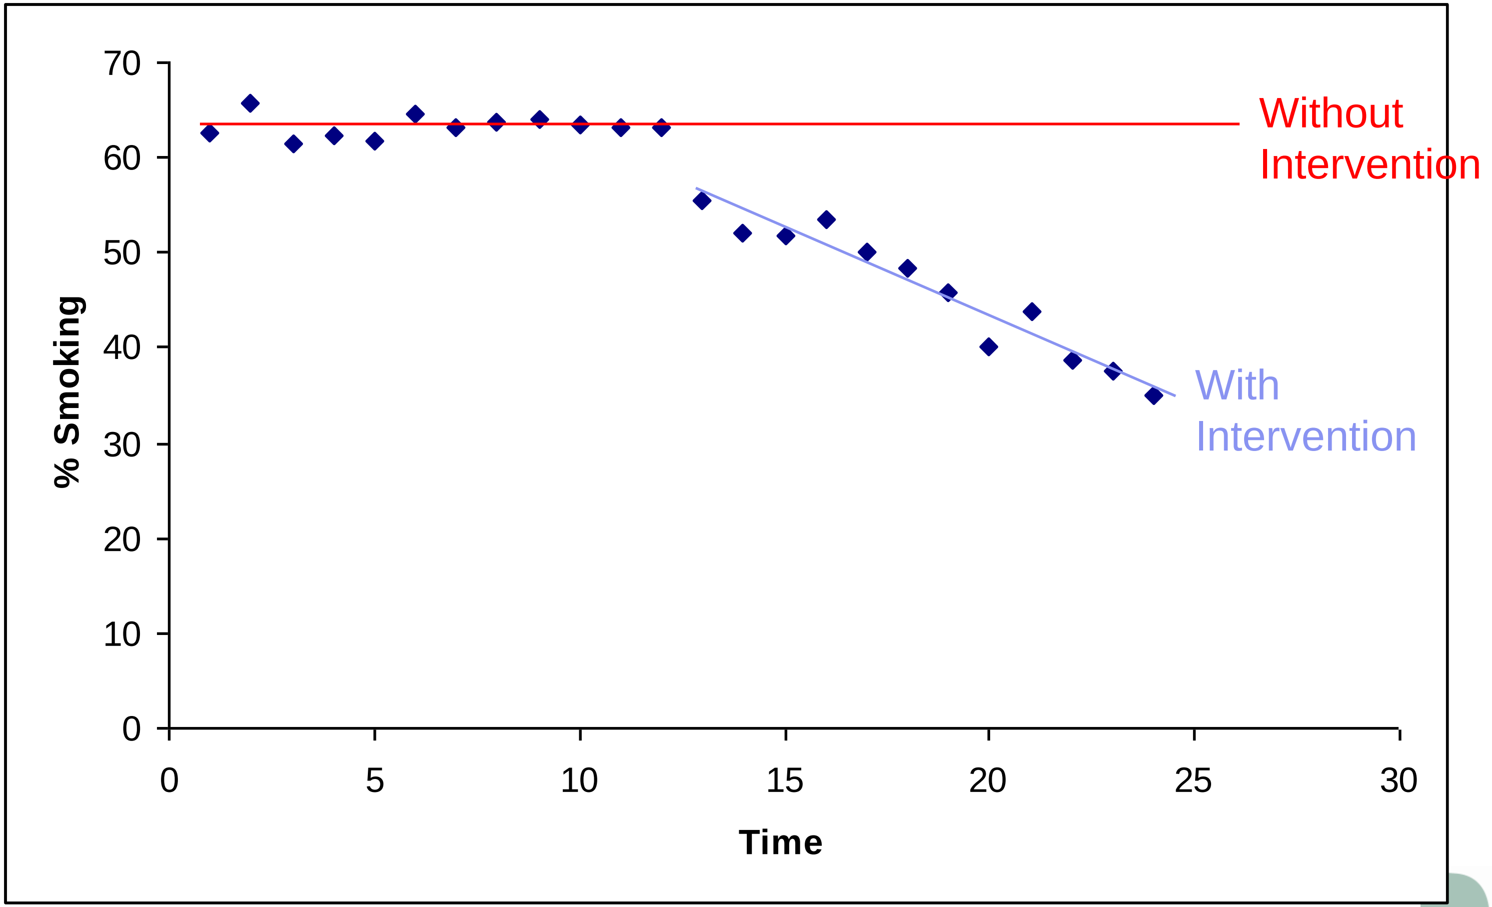

Center for Health Research

Errors in Prescribing

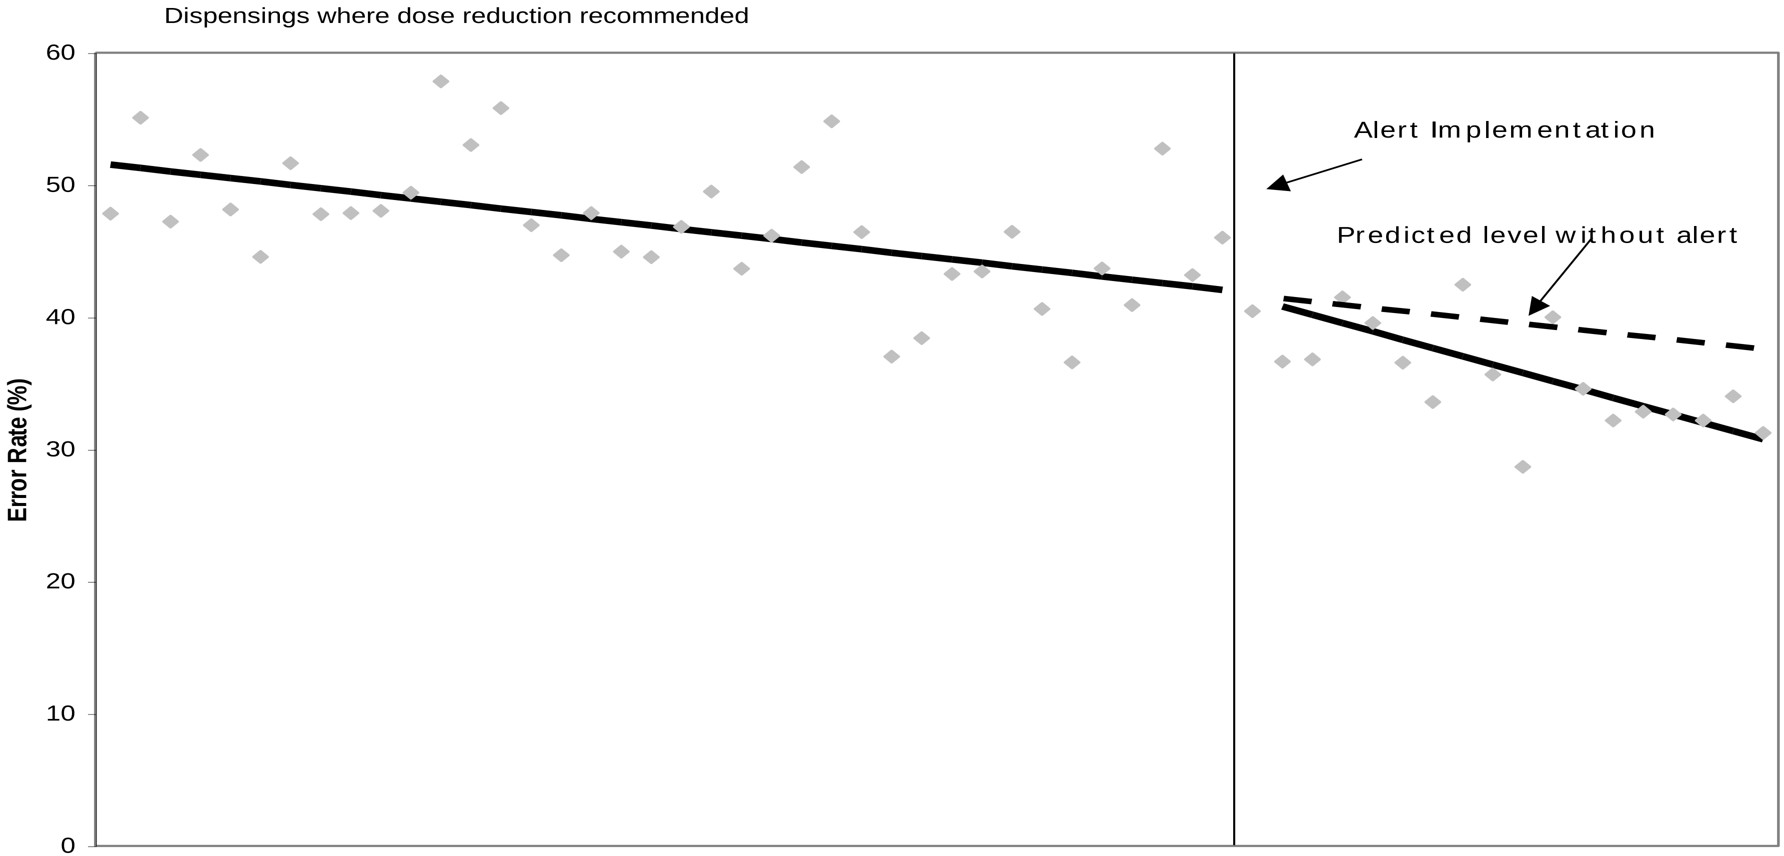

Center for Health Research

Practice Problem

- Percent of 5th graders in a rural county that are able to pass a reading test
- Have historical rates of passing
  - Students tested every six months
  - 15 observations prior to intervention
- Implement a county-wide program to help students read
  - 15 observations post the intervention
- Is the program effective?

Center for Health Research

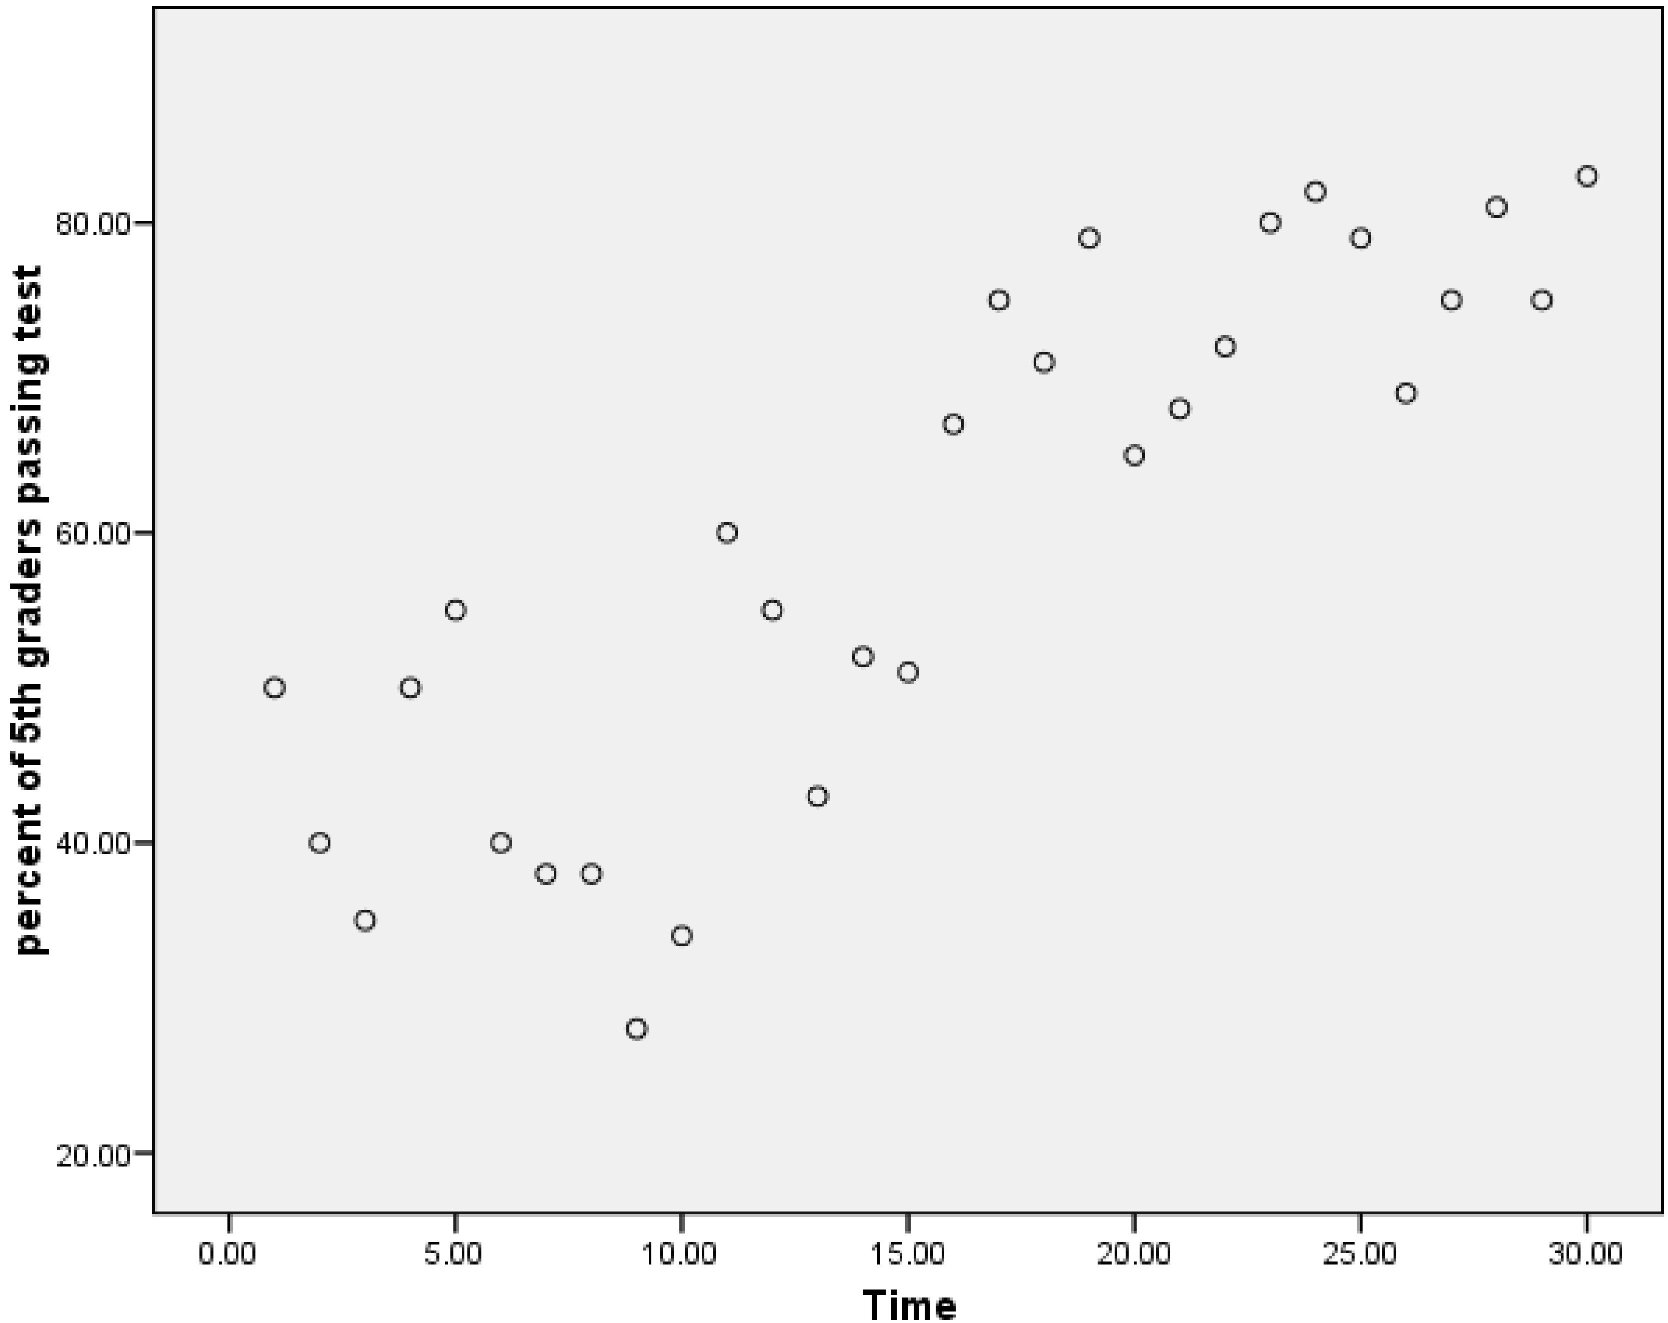
percent of 5th graders passing test
80.00
60.00
40.00
20.00
0.00
5.00
10.00
15.00
20.00
25.00
30.00
Time

What variables do we need in our dataset?

- Reading scores
- Time from the beginning
- Dummy variable for pre or post the reading program implementation
- Time since the program began

Coefficients[a]

Model		Unstandardized Coefficients B	Unstandardized Coefficients Std. Error	Standardized Coefficients Beta	t	Sig.
1	(Constant)	40.829	4.137		9.869	.000
	Time	.471	.455	.242	1.036	.310
	intervention	21.405	5.579	.635	3.837	.001
	postItime	.207	.644	.062	.322	.750

a. Dependent Variable: percent of 5th graders passing test

- What was the percent of 5th graders passing the test at baseline?
  - What about after the intervention?
  - Was the change is level significant?
- What was the slope prior to the intervention?
  - What about after the intervention?
  - Was the change in slope significant?

Center for Health Research

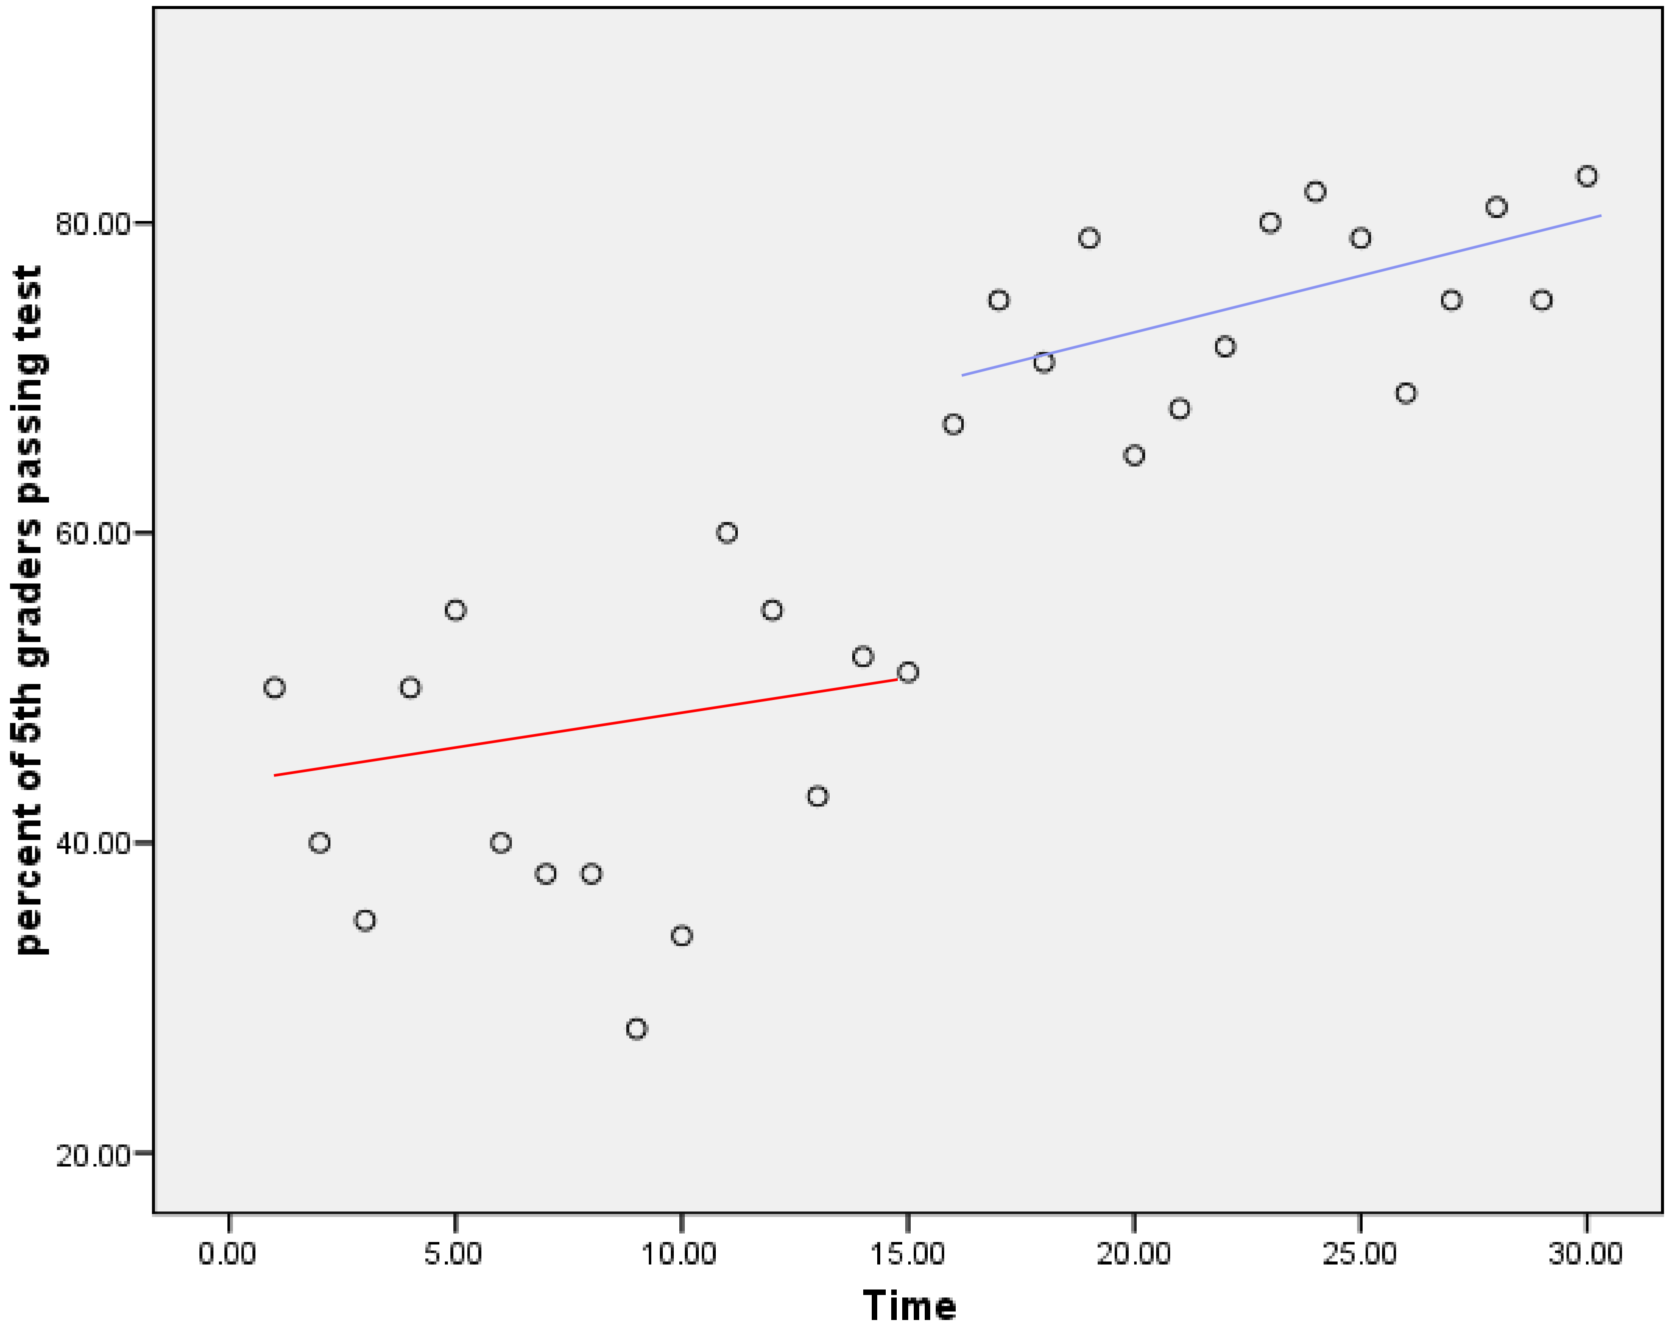

percent of 5th graders passing test
80.00
60.00
40.00
20.00
0.00
5.00
10.00
15.00
20.00
25.00
30.00
Time

Estimating effects at a given time

- What is the percent of kids expected to pass the test at 2 ½ years post the intervention if the intervention had not happened?
  - What is the predicted rate with the intervention?
- Remember time interval is every 6 months
  - 2 ½ years post intervention is post-intervention time of 5

Center for Health Research

Autocorrelation

- Ordinary least squares regression assumes that the error terms associated with each observation (time point) are uncorrelated
- Observations may not be independent
  - Error terms of points close in time can be correlated
- Order
  - First-order autocorrelation – adjacent time points are correlated
  - Second-order autocorrelation – correlation between a time point and 2 months after (or before)
- If don't correct standard errors will be underestimated
  - Significance would be overestimated
- Never fear…if autocorrelation exists it can be controlled for in the analysis

Autoregressive model

- Use plot of residuals over time to assess if autocorrelation exists
  - Random pattern → no autocorrelation
- Durbin-Watson statistics to test for serial correlation of the error terms
  - Nonsignificant means no autocorrelation

Seasonal Trends

- Time of year may influence the rates
  - Screening rates are higher in the fall than any other time of year
  - If a series is stationary (no seasonal trends) the mean is constant over time

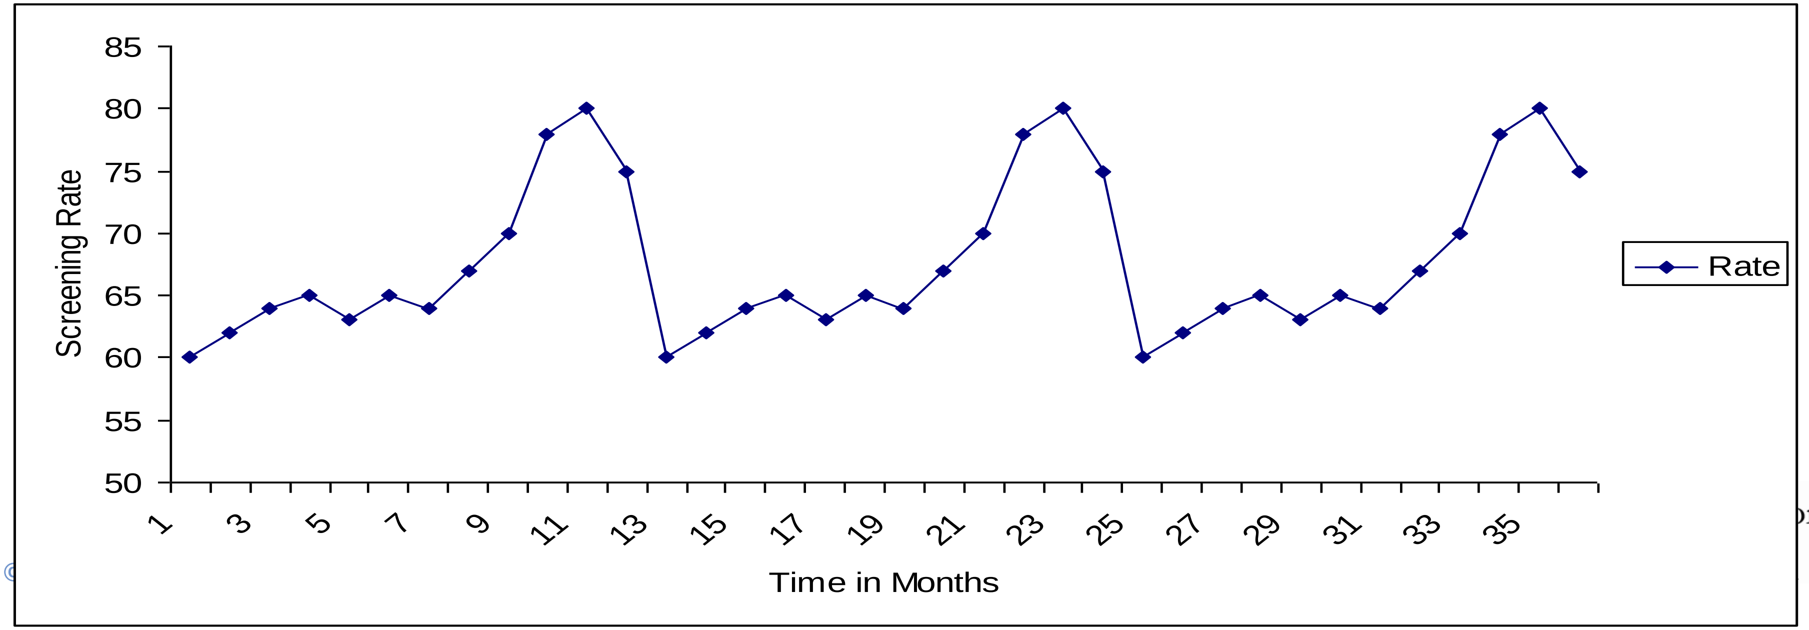

Dealing with seasonality

- Take the difference of the series from one time period to the next
  - Use the differences as the data
  - Might have to take more than one difference
  - Differences can also be lagged
- ARIMA (p, d, q) model
  - p - autoregressive order (number of terms needed)
  - d – non-seasonal differences
  - q – moving average order
- If you are not familiar with these models definitely consult with a statistician
- Interpretation of the model does not change

Center for Health Research

Threats to Internal Validity of Time Series Designs

- History
  - Some other event occurred at the same time as the intervention and so could also explain the pattern of change over time
    - Cigarette machines removed from schools
- Selection
  - The population changed at about the same time as the intervention
    - Parents who favor the program move their kids to that school
- Instrumentation
  - Aspects of the record keeping procedures may change at the same time as the intervention
    - Decreased staff leads to fewer kids being caught smoking

Center for Health Research

Safeguards Against Threats to Internal Validity

- Include a control series for another group
- Examine a different yet similar outcome over the time period that should not be effected by the intervention
- Have multiple implementation time points for the intervention

Center for Health Research

No Treatment Control

- Simultaneously examine another group that would be affected by history, selection, or changes in instrumentation but not the implementation of the intervention

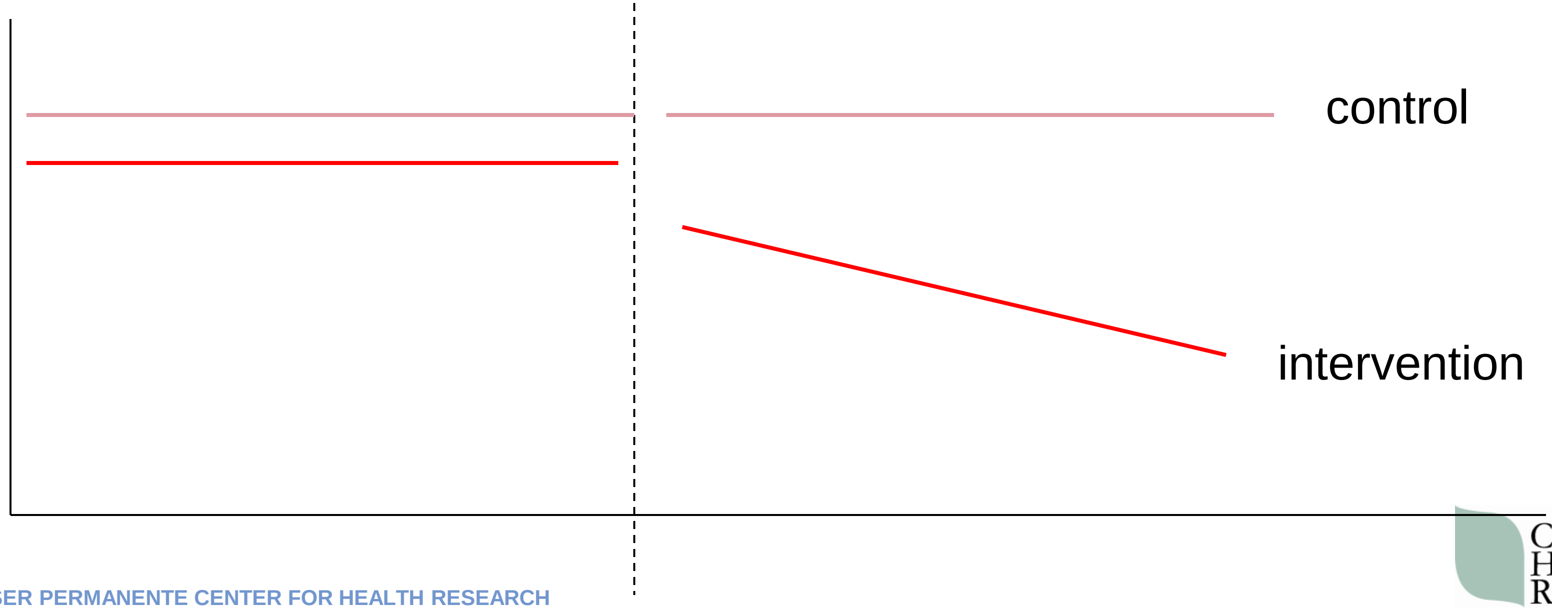

Center for Health Research

Segmented Regression with Two Groups

- Need to incorporate a dummy variable for group and associated interaction terms
  - Group by change in level interaction term
    - Are the control and intervention groups different on the amount of change immediately after the intervention?
  - Group by change in slope interaction term
    - Are the control and intervention groups different on the change in slope from pre to post intervention?

Center for Health Research

Terms in the segmented regression model with two groups

$$Y_t = b_o + b_1 T + b_2 D + b_3 P + b_4 G + b_5 GT + b_6 GD + b_7 GP$$

- G – dummy variable for group
- b_0 – value of DV at baseline
- b_1 - trend prior to intervention
- b_5 - difference between the groups in prior trend
- b_2 – change in level
- b_6 - difference between the groups in change in level
- b_3 – change in trend post intervention
- b_7 - difference between the groups in change in trend

Center for Health Research

Note:

2 obs per time point

Obs	%smoking	T	D	P	G	GT	GD	GP
1	62.5	1	0	0	0	0	0	0
2	65.8	1	0	0	1	1	0	0
3	61.4	2	0	0	0	0	0	0
4	62.3	2	0	0	1	2	0	0
5	61.8	3	0	0	0	0	0	0
6	64.6	3	0	0	1	3	0	0
7	63.1	4	0	0	0	0	0	0
8	63.8	4	0	0	1	4	0	0
9	64.1	5	0	0	0	0	0	0
10	63.5	5	0	0	1	5	0	0
11	63.1	6	0	0	0	0	0	0
12	63.3	6	0	0	1	6	0	0
13	55.4	7	1	1	0	0	0	0
14	52.1	7	1	1	1	7	1	1
15	51.7	8	1	2	0	0	0	0
16	53.4	8	1	2	1	8	1	2
17	50.1	9	1	3	0	0	0	0
18	48.3	9	1	3	1	9	1	3
19	45.9	10	1	4	0	0	0	0

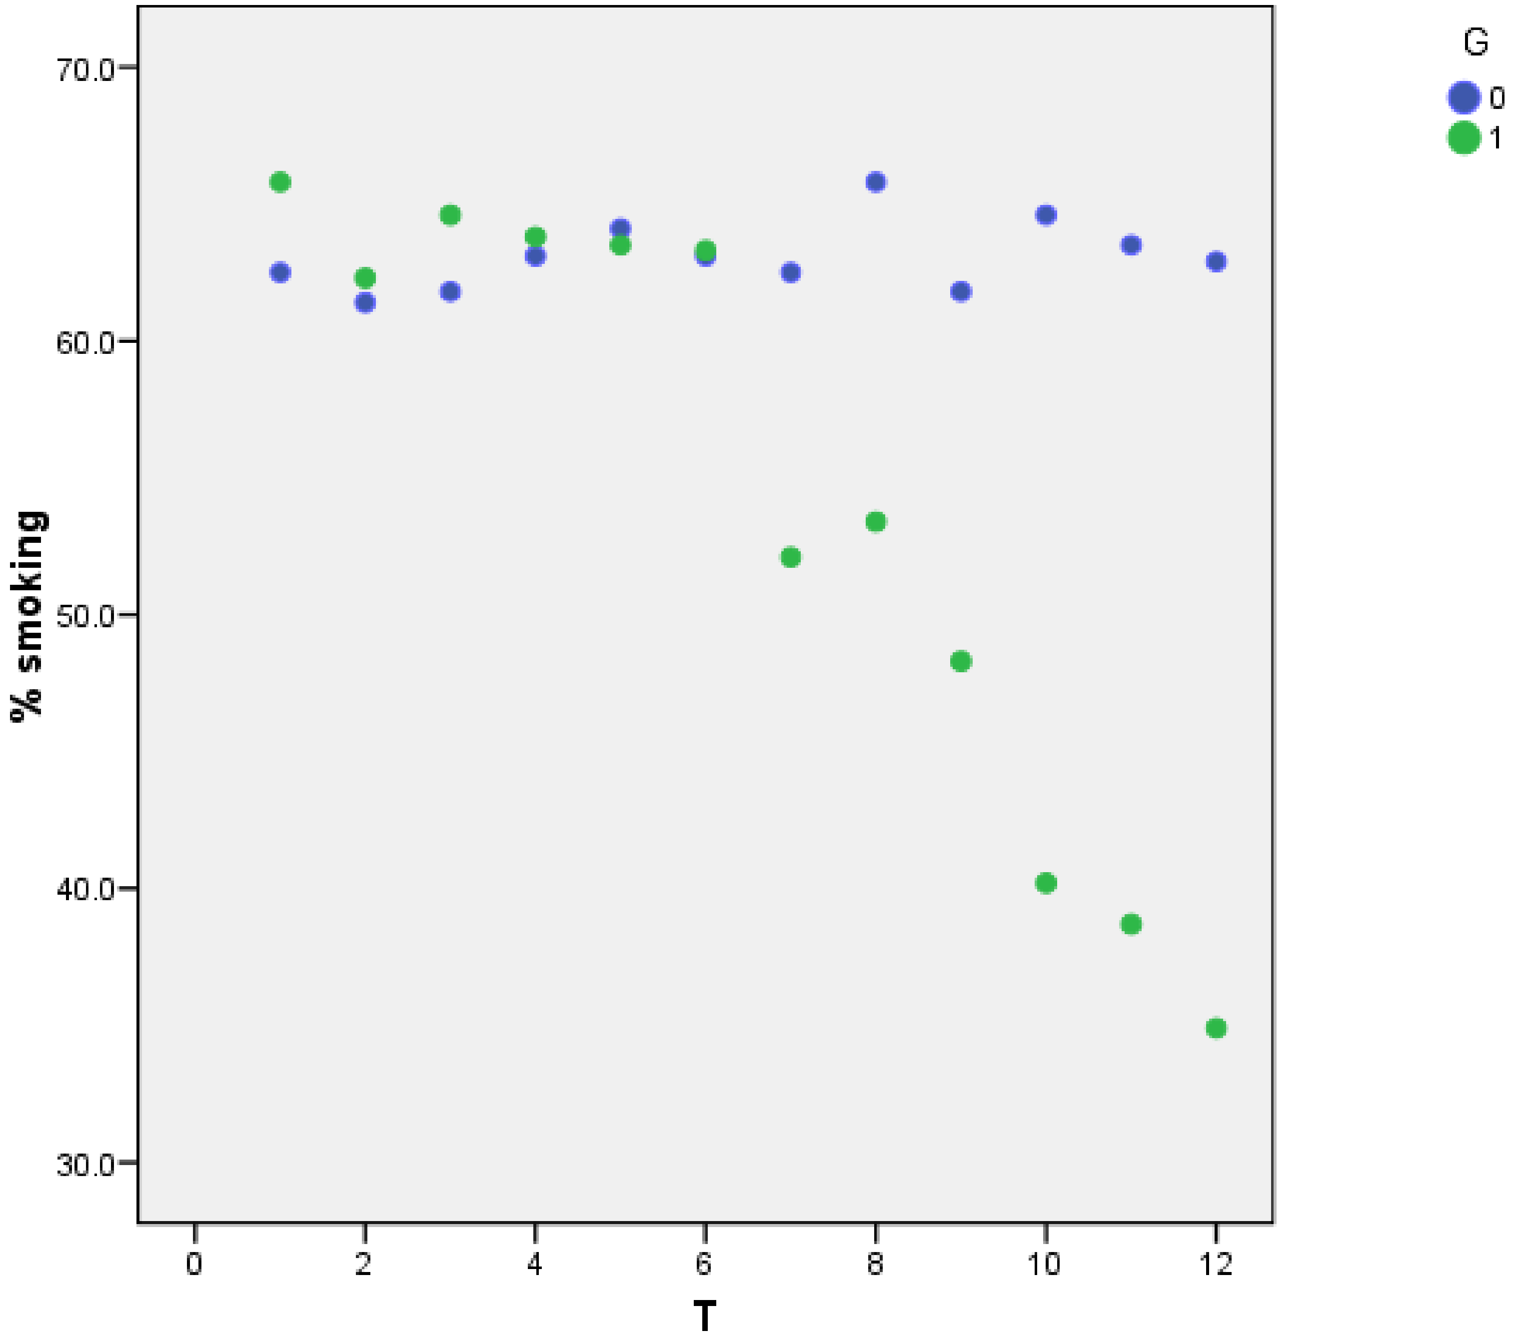
G
0
1
70.0
60.0
50.0
40.0
30.0
% smoking
0
2
4
6
8
10
12
T

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	61.427	1.507		40.755	.000
	T	.354	.387	.137	.915	.374
	D	.174	1.909	.010	.091	.928
	P	-.414	.547	-.099	-.757	.460
	G	3.427	2.132	.192	1.608	.127
	GxT	-.631	.547	-.288	-1.154	.266
	GxD	-4.945	2.700	-.240	-1.831	.086
	GxP	-3.257	.774	-.636	-4.208	.001

a. Dependent Variable: % smoking

There was a significant difference between the control and intervention groups in the change in slope from pre to post intervention (GxP, p<.001).

The immediate change after the intervention was not significantly different for the control and intervention groups (GxD, p=.086)

Center for Health Research

Other Control Series

- Simultaneously examine another outcome that would be affected by history, selection, and instrumentation but not the implementation of the intervention

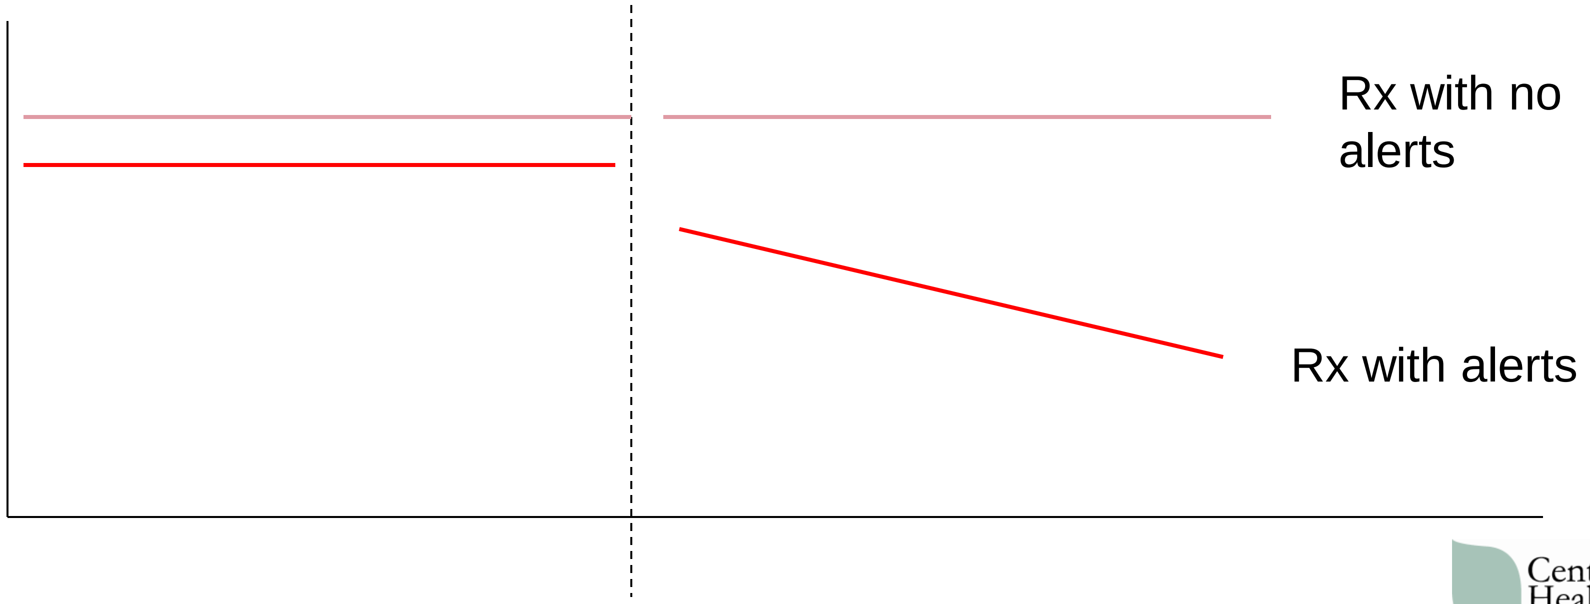

Center for Health Research

Other Control Series

- The dependent variables are different so they cannot be incorporated into the same regression model
- Ideal is to see a significant level and slope change in the outcome of interest but not in the control variable

Center for Health Research

Time Series Example:
Safety in Prescribing (SIP)

- David Smith - Center for Health Research
- Effect of patient-specific decision support tool in the EMR on reducing risky prescribing
- Outcome rate of prescribing
  - Medications that should not be prescribed to the elderly

- Smith, Perrin, Feldstein, Yany, Kuang, Simon, Sittig, Platt, Soumerai. The impact of prescribing safety alerts for elderly persons in an electonic medical record: An Interrupted time series evaluation. *Arch Intern Med. 2006;166:1098-1104*.
- Feldstein, Smith, Perrin, Yang, Simon, Krall, Sittig, Ditmer, Platt, Soumerai. Reducing warfarin medication interactions: An interrupted time series evaluation. *Arch Intern Med. 2006;166:1009-1015*.

Funded by AHRQ U18HS11843

Center for Health Research

Alerts for the Elderly

- Time Series
  - 13 points prior to the alerts, 26 points post alert implementation
- Dependent Variable
  - Prescribing rates per 10,000
- Target and control groups
  - 65+ years old - target
  - Under 65 years old - control
    - Alerts fired regardless of age of the patient
- Two types of antidepressants
  - Anitriptyline (non-preferred for elderly)
  - Nortriptyline (preferred for elderly)
- Used autoregressive models
  - Seasonality in depression

Center for Health Research

Time series of nonpreferred (A) and preferred (B) agent dispensings per 10 000 by age

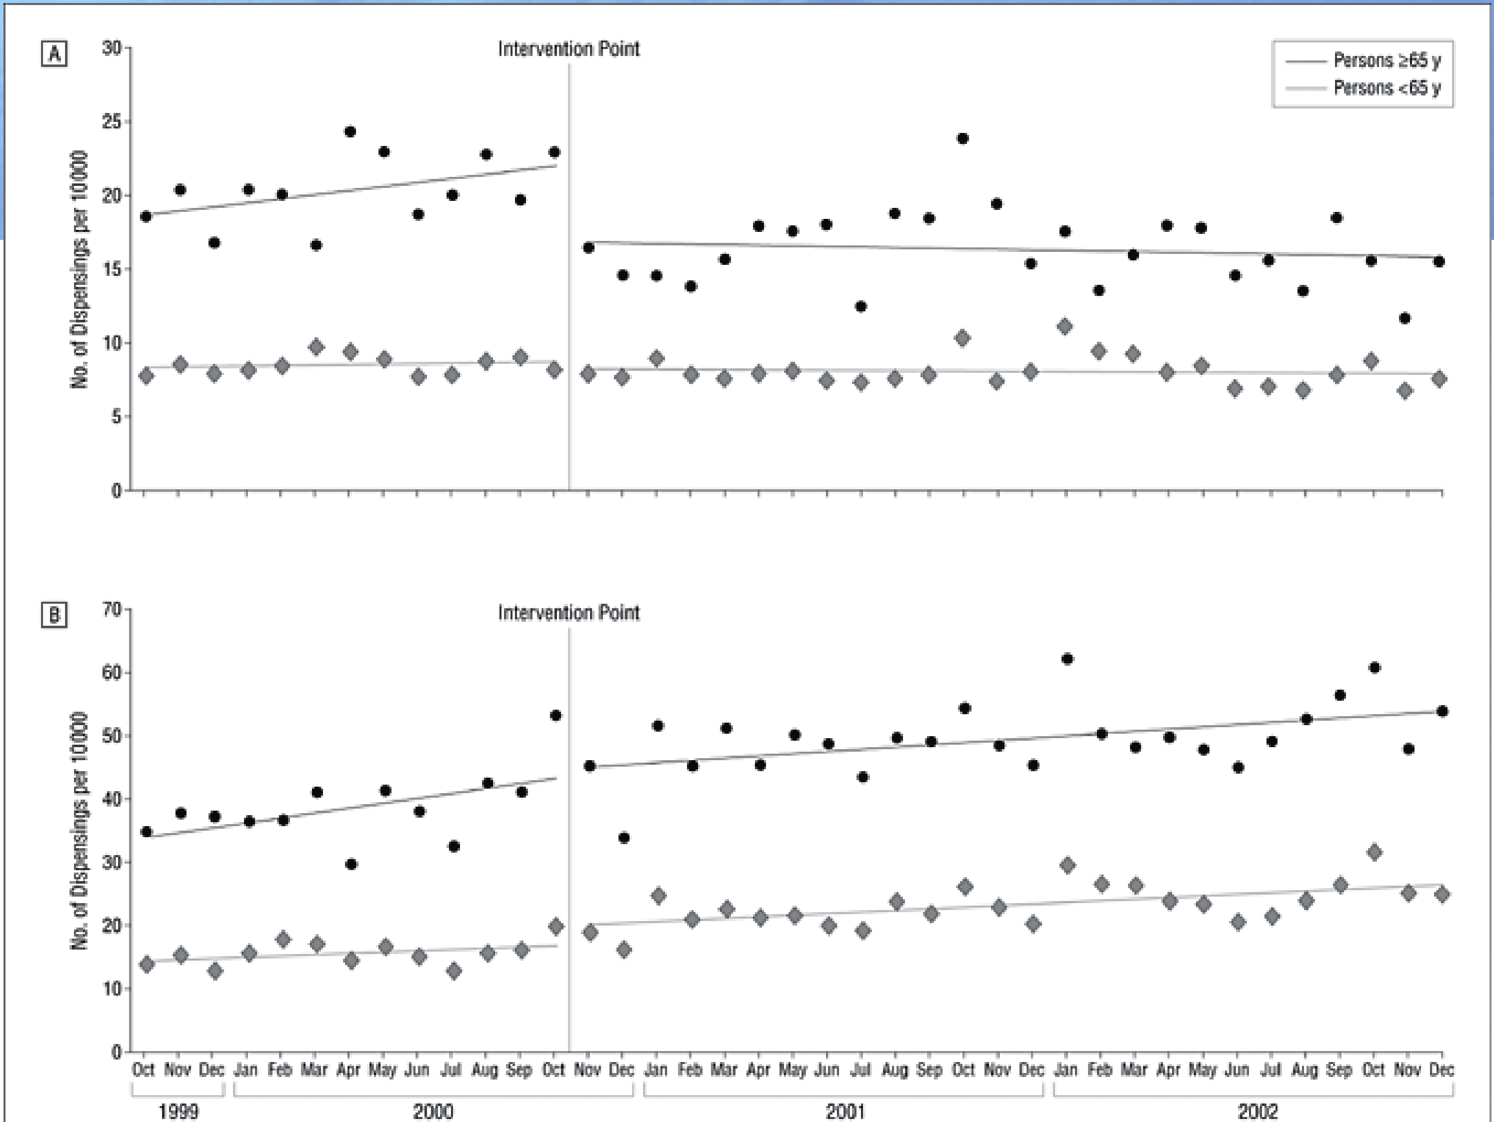

Smith, D. H. et al. Arch Intern Med 2006;166:1098-1104.

ARCHIVES OF INTERNAL MEDICINE

Multiple Implementation Periods

- If the intervention is phased in can compare the timing of the interruption in the time series

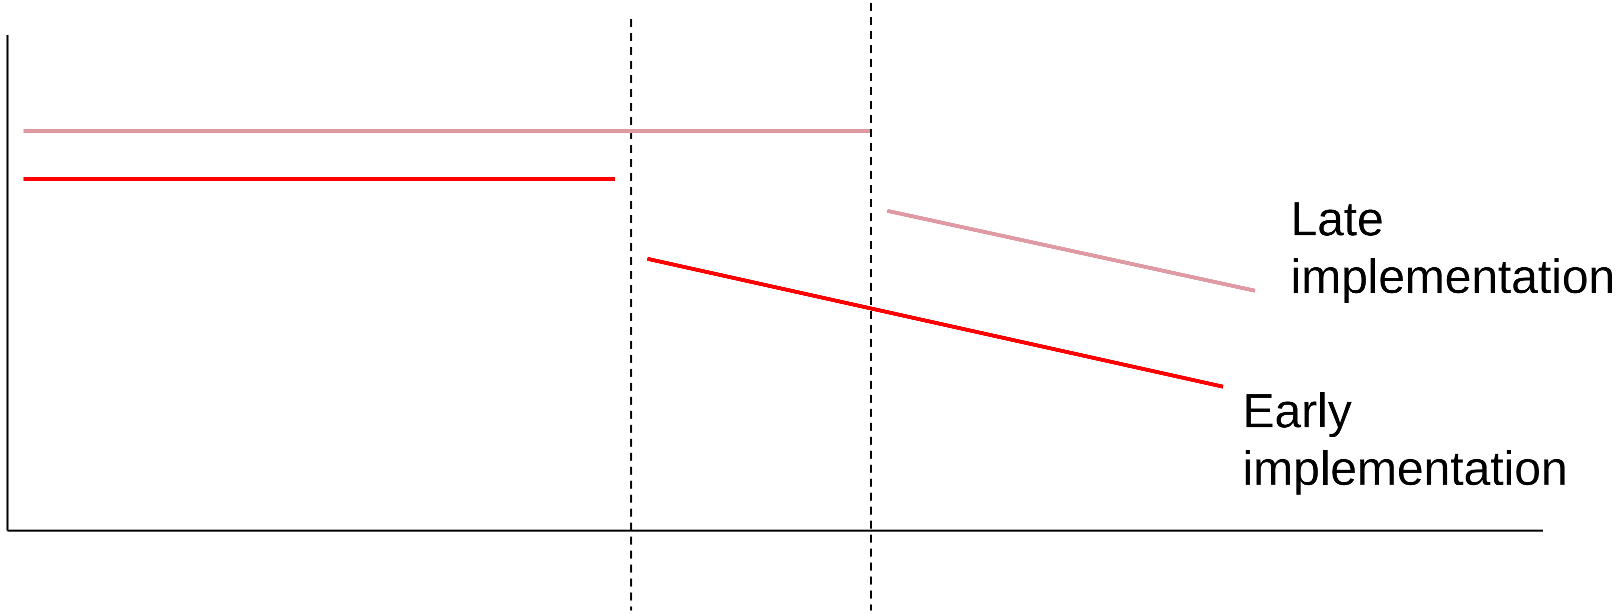

Center for
Health
Research

Segmented Regression with Multiple Time Periods

$$Y_t = b_o + b_1 T + b_2 D_A + b_3 P_A + b_4 D_B + b_5 P_B$$

- T is time from baseline
- D_A is a dummy for the first implementation
  - Coded 0 for prior to 1st implementation, 1 after
- P_A is time since first implementation
- D_B is a dummy for the second implementation
  - Coded 0 for prior to 2nd implementation, 1 after
- P_B is time since second implementation

Center for Health Research

What would the data set look like for 2 implementation time points?

- Time since baseline
- Dummy for first intervention
- Time since first intervention
- Dummy for second intervention
- Time since second intervention

Center for Health Research

Obs	%smoking	T	DA	PA	DB	PB
1	62.5	1	0	0	0	0
2	65.8	2	0	0	0	0
3	61.4	3	0	0	0	0
4	62.3	4	0	0	0	0
5	61.8	5	0	0	0	0
6	64.6	6	0	0	0	0
7	63.1	7	1	1	0	0
8	63.8	8	1	2	0	0
9	64.1	9	1	3	0	0
10	63.5	10	1	4	0	0
11	63.1	11	1	5	0	0
12	63.3	12	1	6	0	0
13	55.4	13	1	7	1	1
14	52.1	14	1	8	1	2
15	51.7	15	1	9	1	3
16	53.4	16	1	10	1	4
17	50.1	17	1	11	1	5
18	48.3	18	1	12	1	6
19	45.9	19	1	13	1	7
20	40.2	20	1	14	1	8

$Y_t = b_o + b_1 T + b_2 DA + b_3 PA + b_4 DB + b_5 PB$

- b_0 – value of dependent variable at baseline
- b_1 - trend prior to 1st intervention implementation
- b_2 - difference between the last point prior and first point post the 1st intervention implementation
- b_3 – change in trend from pre to post 1st implementation
- b_4 - difference between last point in 1st implementation period and first point in 2nd implementation period
- b_5 – change in trend from 1st to 2nd implementation periods

Center for Health Research

Multiple Implementation Periods

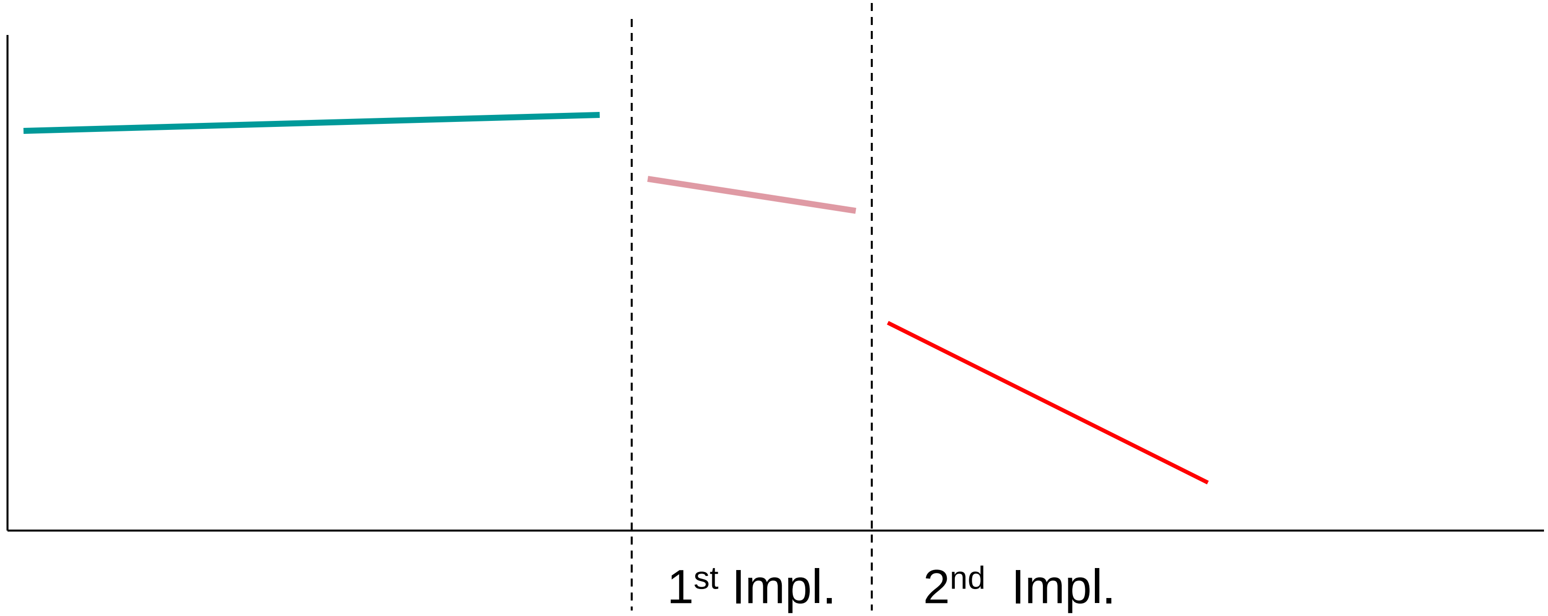

Center for Health Research

Statistical Power Issues

- Power is greatest when the intervention occurs in the middle of the time series
- Stable baseline period will increase power to detect a change
  - Following same population across time reduces variability in the time series
  - Basing each rate on a large sample size will reduce variability
    - Might consider aggregating (quarterly instead of monthly)
- Studies with moderate effect sizes have been significant with 12 points pre and post

Center for Health Research

Delayed Effects

- Time series is strongest when the intervention produces an immediate effect
  - Intervention may take time to affect the outcome
  - Intervention effect could be reduced if it is implemented slowly over time
    - Include the implementation period as a separate time period or
    - Consider taking that time period out of the analysis
    - Need to know the length of the implementation period
    - Collect supplementary data about the diffusion of the intervention over time

Center for Health Research

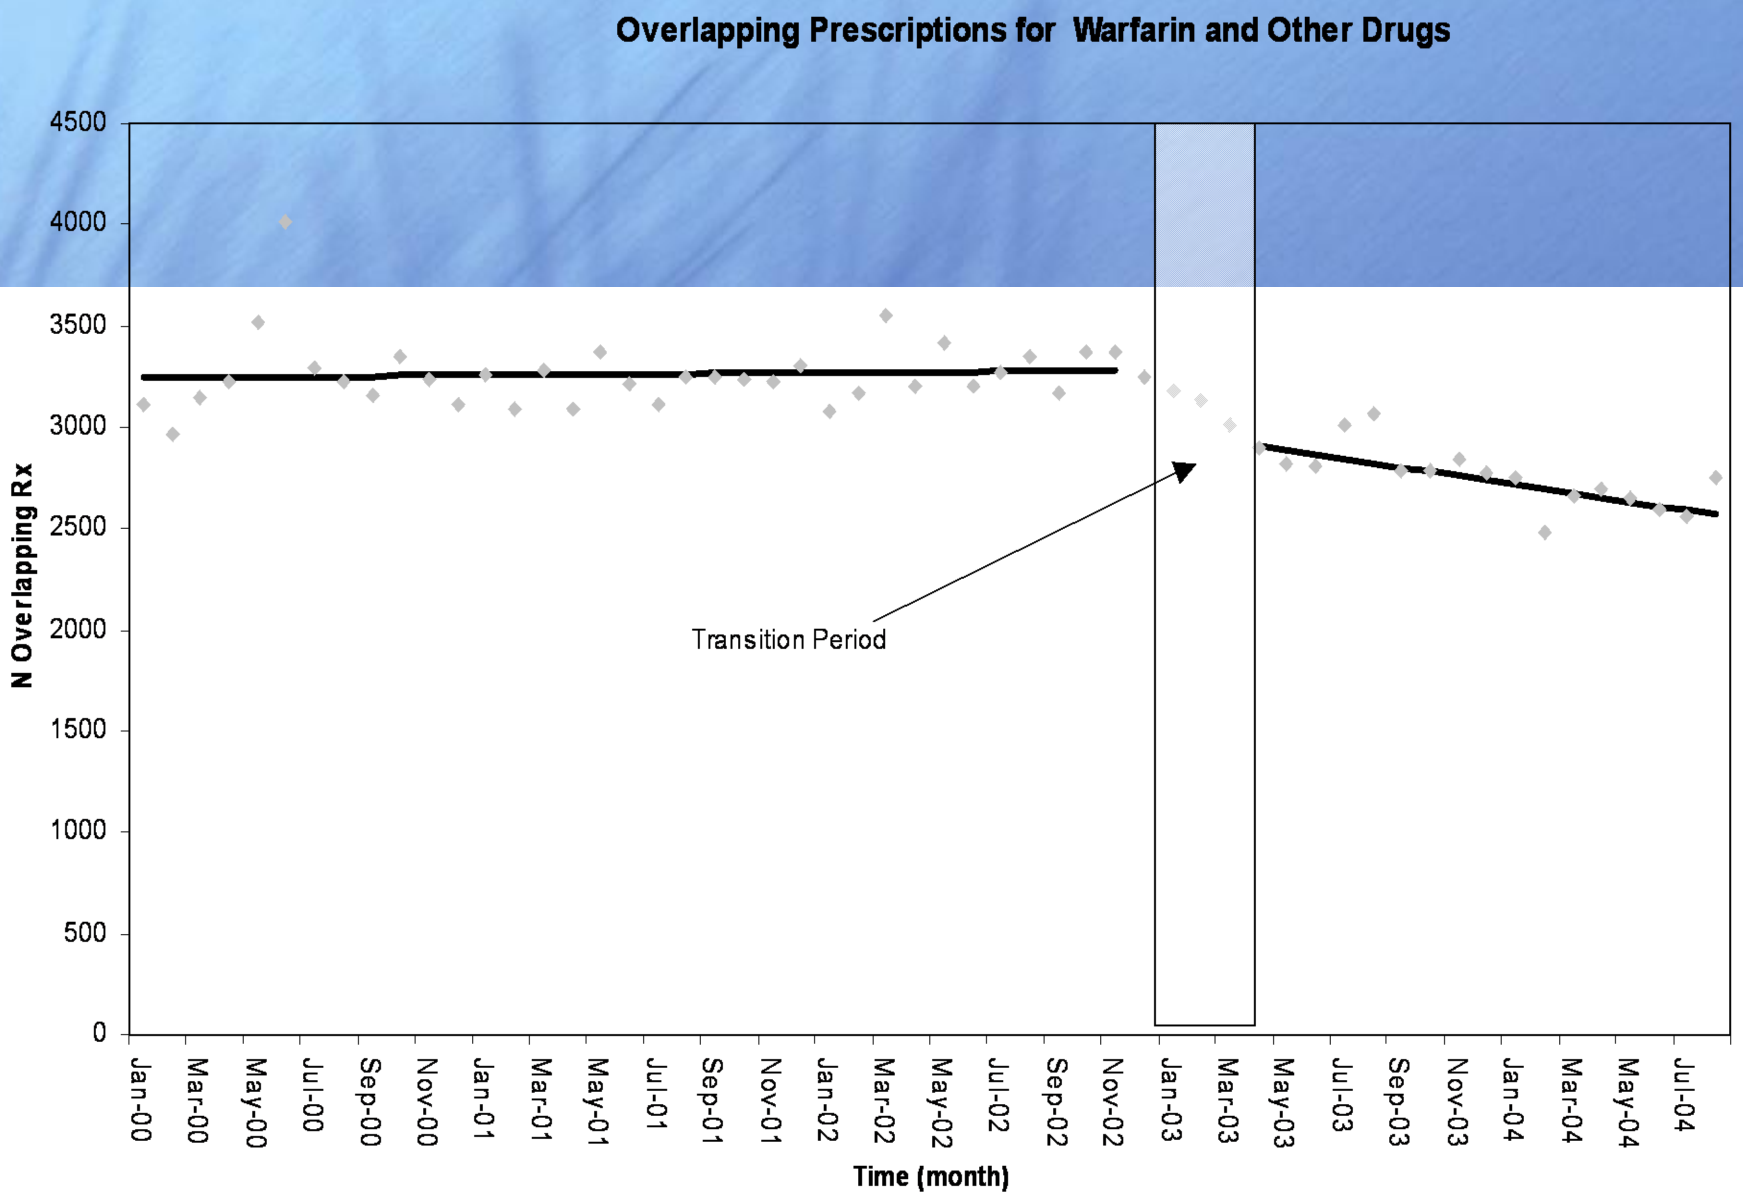

Feldstein, A. C. et al. Arch Intern Med 2006;166:1009-1015.

Center for Health Research

Summary

- Time series designs are a strong quasi-experimental alternative to randomization
- External validity of the design may be high as it occurs in a natural setting
- Segmented regression is a robust technique to test immediate and sustained changes associated with the intervention
  - While controlling for secular trends
- The time series design can be strengthened with control series or alternative outcomes that should not be affected by the intervention

Center for Health Research

Calculating absolute difference and relative change for practice problem

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	40.829	4.137		9.869	.000
	Time	.471	.455	.242	1.036	.310
	intervention	21.405	5.579	.635	3.837	.001
	postItime	.207	.644	.062	.322	.750

a. Dependent Variable: percent of 5th graders passing test

Y = 40.83 + (.47*20) + (21.41*1) + (.21*5) = 72.69% passing test intervention

Y = 40.83 + (.47*20) + (21.41*0) + (.21*0) = 50.23% passing without the intervention

22.46% (72.69-50.23) absolute difference

44.71% (22.46/50.23*100) relative percent change

Center for Health Research