

Lecture 4

Let's open to today with a central economic concept – [opportunity cost](#).

The opportunity cost of a particular course of action is the value of the next best alternative course of action that could have been chosen.



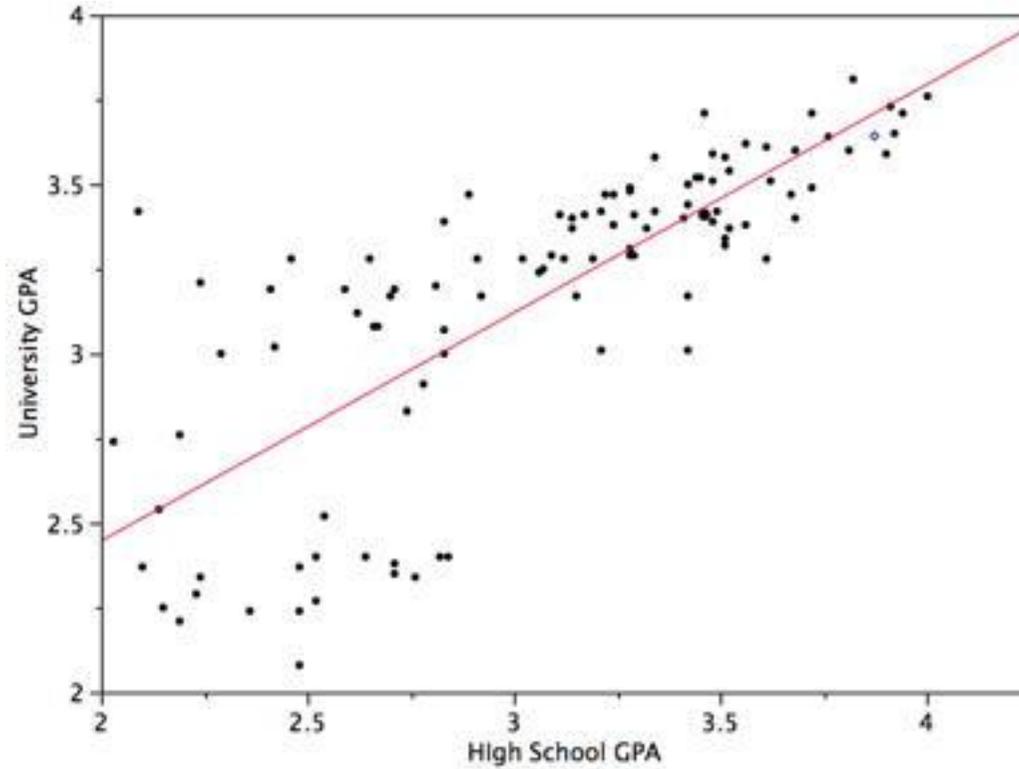
Imagine you are out for a hike. The opportunity cost of going left is the pleasure you would have received if you had gone right instead. The opportunity cost of going right is the pleasure you would have received from going left.

Time is a scarce commodity. You cannot simultaneously go right and left. It is important to realize that your choice is not free even though you do not pay any money.



IF YOU SPEND NINE MINUTES OF YOUR TIME TO SAVE A DOLLAR, YOU'RE WORKING FOR LESS THAN MINIMUM WAGE.

Next let's recall quickly the basic ideas behind [linear regression](#).



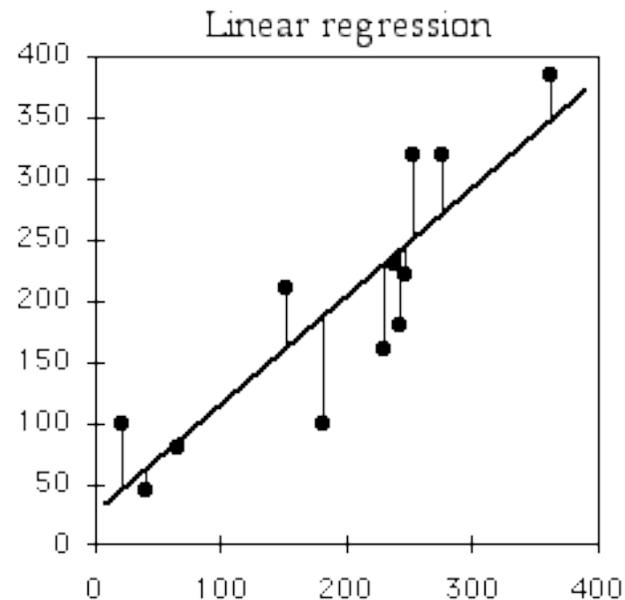
(“GPA” stands for “Grade Point Average” which is a measure of how well students perform in school.)

Here are some general observations about the picture on slide 4.

1. The red line is meant to fit the points in the scatterplot as well as possible and is known as the *linear regression line*.
2. The positive slope of the red line suggests that higher grades in high school are associated with higher grades in university.
3. It is not clear that there is a [causal relationship](#) between high school GPA and university GPA. It could be, for example, that higher IQ causes higher GPA in both high school and university. More generally, there can be a third factor driving GPA at both levels.

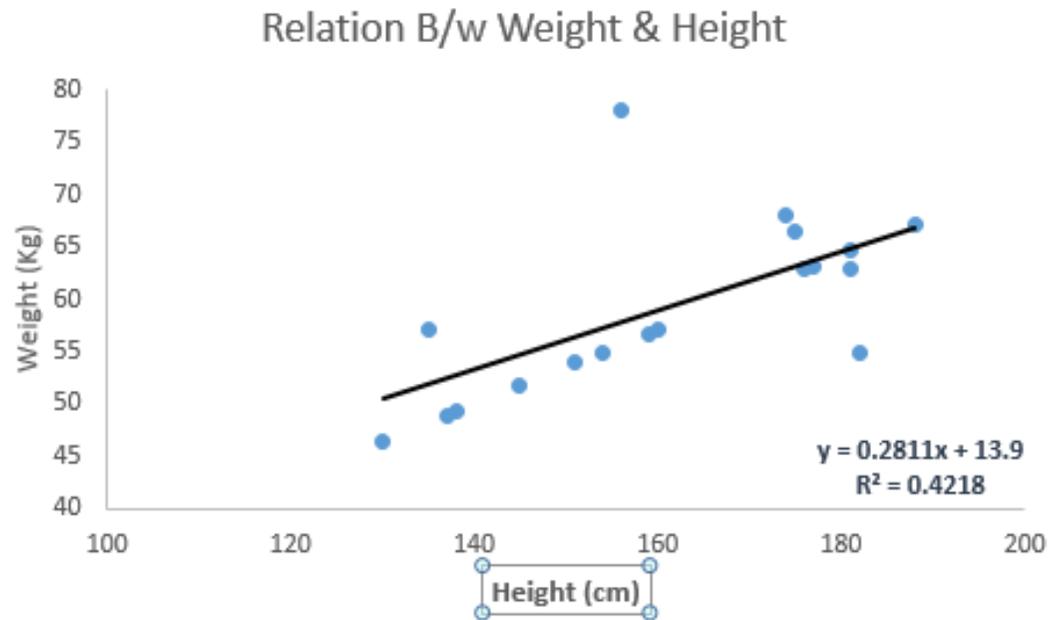
Here's a useful thought experiment for thinking about causality in this case. Imagine your old high school contacts you and says that they just detected some errors in your records and they are raising some of your marks. Would this cause you to suddenly perform better at RHUL? Probably not, which suggests that it is not the high school marks themselves that drive your university marks.

In what sense is a [linear regression line the best fit](#)? Consider this simple picture:



Take the distances shown, square them and add these numbers up. You get a measure of how well the line in the picture fits the data points. If I draw a different line you could then measure the distances of all the same points to this line, square them and add them up. *The linear regression line is the particular line that minimizes the sum of the squared distances to the data points.* Software like [Stata](#) or [Excel](#) will find these linear regression lines for you.

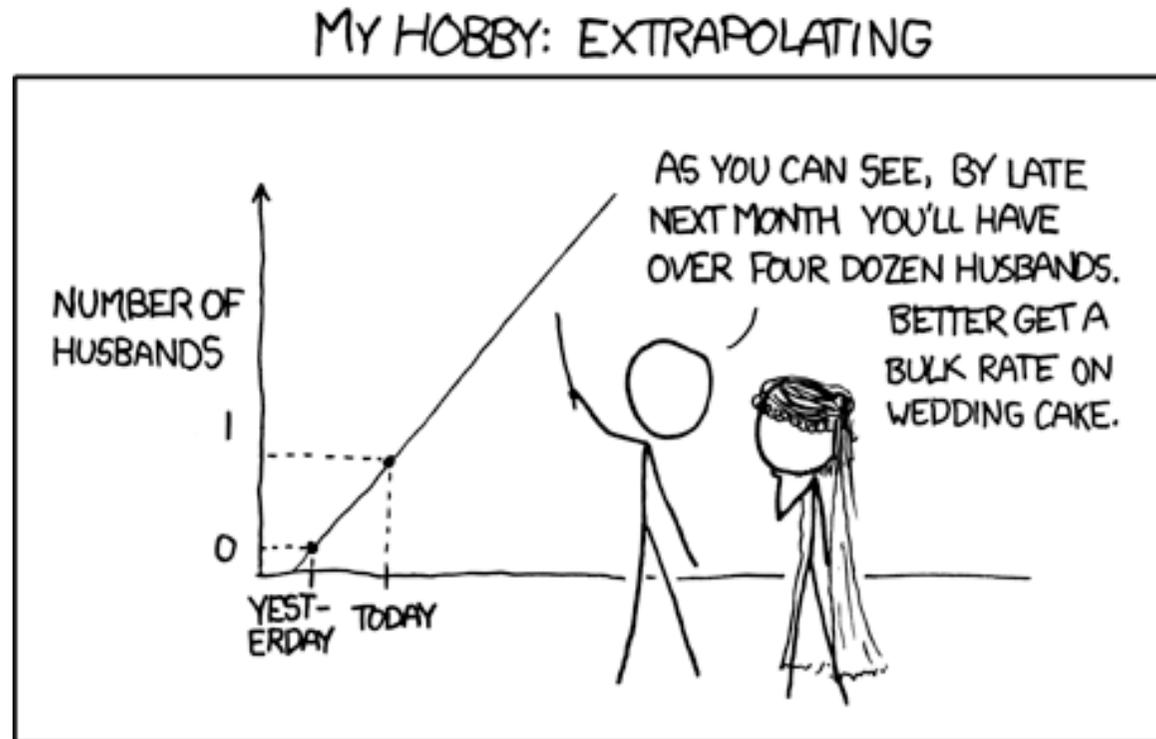
The picture below gives not only a linear regression line but also the equation of this line.



The interpretation is that 1 more centimetre of height is associated, on average, with an additional 0.28 kilograms of weight.

We can also use this equation to predict that a person who is 160 centimetres high will weigh around 58.7 kilograms. Of course, this prediction could turn out to be way off.

Notice that it is hazardous to make predictions outside of the range for which you have no data. For example, we can predict that a person who is 100 centimetres tall will weight 41.9 kilograms but we should be cautious about this prediction since we have no observations for anyone below 125 centimetres tall. It could be that this is a good prediction but it assumes that the straight line extends beyond the region for which we have decent evidence that it works well.



To rebel or not to rebel, that is the question.

This week we have two Berman et al. papers so we need to distinguish between them.

Berman et al. 1 will be “[Do Working Men Rebel? Insurgency and Unemployment in Afghanistan, Iraq, and the Philippines?](#)”

Berman et al. 2 will be “[Modest, Secure and Informed: Successful Development in Conflict Zones.](#)”

Berman et al. 1 studies the relationship between insurgent activity and unemployment.

Many people assume they understand how this relationship works. For example, here is a [quote from the top U.S. Army Commander in Iraq back in 2006](#):

“...employment is absolutely critical to what we're doing to take the angry young men off the street.”

We can characterize this view as the *opportunity cost* theory of insurgency. The idea is that people considering joining an insurgency think about what their next best opportunity is. This is likely to be the wage they would earn in the legal economy if they do not join.

If the unemployment rate is high then the next best opportunity to joining an insurgency could be unemployment, in which case joining up may look like a pretty good option.

This opportunity cost view is sensible and intuitive but a little reflection (or else just reading the Berman et al. paper....) indicates that there are other mechanisms that can create either positive or negative correlations between unemployment and violence. The following three mechanisms go against the simple opportunity cost story.

1. Security measures should reduce violence but may also reduce economic activity, thereby raising the unemployment rate.

For example, in 2007 Baghdad erected barriers between Sunni and Shiite neighbourhoods, a measure that probably contributed to reducing sectarian killings. But these barriers also impeded the free flow of goods and labour around the city, with a negative impact on the economy.

There appears to be a similar phenomenon at work when border crossings between Israel, the West Bank and Gaza Strip are closed – less violence but higher unemployment.

2. A high unemployment rate can help counterinsurgent forces because it becomes easier to purchase information about the insurgency from unemployed people holding such information.

In other words, some people may be reluctant to provide information to the army but if they are unemployed and desperately trying to support their families they may do so for money. The army may be able to use these tips to capture insurgents and reduce violence.

3. Some violence could be aimed at stealing economic resources. When the economy is doing well there are more resources to steal, possibly provoking violence as have-nots try to steal from haves. When the economy is doing poorly there may be less such violence since there are not so many resources available for stealing.

Of course, when the economy does well the unemployment rate also tends to be low.

Together, these two considerations suggest a *negative relationship* between unemployment and insurgent violence.

All the channels in slides 11-12 push in the direction of a *negative relationship between unemployment and violence*, that is, higher unemployment is associated with lower violence and vice versa in these stories.

Berman et al. 1 also offers a further reason (in addition to the opportunity cost theory) for a *positive* relationship between unemployment and violence: violence itself may disrupt the economy, causing the unemployment rate to increase.

It is important to recognize that several of the above channels can operate simultaneously, making it extremely difficult to determine causal effects.

If violence causes unemployment and unemployment also causes violence then we have what is known as a *reverse causation* problem. I will return to this issue later in the lecture.

The basic idea of the Berman et al. 1 paper is to estimate an equation like the following:

$$V_r = \alpha + \beta u_r + \varepsilon_r$$

where V stands for “violence”, r is “region,” α and B are estimated coefficients and ε is a random shock that hits at the region level. You can think of the shocks as explaining why the points in a scatterplot do not fall exactly on a regression line.

The opportunity cost theory suggests that the coefficient estimated for B should be positive – more unemployed “angry young men” should cause more violence.

The authors use data from three different wars: Afghanistan, Iraq and the Philippines. The following table summarizes the data.

Table I. Summary Statistics for Population, Unemployment, and Violence

	Variable	Mean	Std. Dev.	Min	Max
	Population (district)	605,340	464,106	10,966	1,624,058
	Unemployment (rate)	0.117	0.063	0	0.495
	Sunni vote share (governorate)	0.207	0.250	0	0.917
	Population density (1,000/km ²)	1.53	3.13	.000286	13.62
	SIGACT incidents/1,000	0.258	0.504	0	8.54
Iraq (district/quarter)	Iraq Body Count (IBC) incidents/1,000				
	Insurgent	0.005	0.011	0	0.240
	Sectarian	0.013	0.026	0	0.340
	Coalition	0.002	0.006	0	0.140
N = 312	IBC civilian casualties/1,000				
	Insurgent	0.018	0.043	0	0.611
	Sectarian	0.046	0.159	0	1.70
	Coalition	0.015	0.102	0	1.32
	Insurgent precision	0.923	0.159	0	1

(continued)

Table I. (continued)

	Variable	Mean	Std. Dev.	Min	Max
Philippines (province/year) N = 624	Population (province)	1,002,281	1,271,405	15,095	11,321,875
	Unemployment (rate)	0.105	0.039	0	0.201
	Percent Muslim	0.053	0.181	0	0.993
	Rebel initiated incidents/1,000	0.010	0.024	0	0.213
	All noncriminal incidents/1,000	0.034	0.070	0	1.13
	Civilian casualties/1,000	0.015	0.042	0	0.659
	Insurgent precision	0.816	0.309	0	1
Afghanistan (district/month) N = 2,160	Population (district)	77,7804	1,410,510	2,462	4,017,898
	Unemployment (rate)	0.241	0.206	0	1
	Percent Pashtun	0.385	0.389	0	1
	Percent males	0.522	0.094	0	1
	Average household size	9.56	1.73	5.84	21.1
	AGE incidents/1,000	0.049	0.152	0	2.18
	JOIIS incidents/1,000	0.083	0.225	0	2.90
	CIDNE incidents/1,000	0.046	0.154	0	2.13

Source: Unemployment and population: Iraq Living Conditions Survey (ILCS; March and April 2004), World Food Program Food Security and Vulnerability Analysis in Iraq (June and July 2005; November and December 2007), Republic of the Philippines Census Organization, Labor Force Survey (LFS) for 1997 through 2003 and 2006, Afghan National Quarterly Assessment Report (ANQAR) surveys wave 1 (September 2008) through wave 5 (September 2009), LandScan™ population. Violence: Empirical Studies of Conflict Philippine project, internal security incidents, Armed Forces of the Philippines, 1975–2008; Multi-National Forces Iraq SIGACTS III Database; Joint Operations—INTEL Information System (JOIIS), Combined Information Data Network Exchange (CIDNE). Iraq Body Count. Other: Sunni vote share in Iraq is based on the December 2005 elections. Muslim population share in the Philippines is from the 1995, 2000, and 2007 censuses. Pashtun population share, male population share, adult population without any formal schooling share, and average household size in Afghanistan are from ANQAR. Note: All means are population weighted.

The actual equations estimated are slightly more complicated than the basic one on slide 15 because they incorporate a time dimension plus various [controls](#).

$$V_{r,t} = \alpha_r + \beta u_{r,t} + \gamma_t + \textit{controls} + \varepsilon_{r,t}$$

where t represents time which is measured in quarters, months or years depending on the regression and the γ_t are other estimated parameters.

The γ_t coefficients take care of the problem that, for example, violence and unemployment may both be increasing over time even though there is no causal relationship between the two.

Slide 18 gives the main table from the paper.

Table 2. Unemployment and Violent Incidents in Iraq, the Philippines, and Afghanistan

		Linear Regression			
DV Region		Incidents /1,000 All	Incidents /1,000 Baghdad	Incidents /1,000 All	Incidents /1,000 Baghdad
Iraq (district/quarter)	Unemployment	-1.118*** (0.542)	-4.254*** (1.460)	-0.709* (0.552)	-1.847*** (0.566)
	Observations	312	27	312	27
	R ²	0.23	0.35	0.06	0.16
	Controls	Ethnicity	Ethnicity	District FE	District FE
	DV Region	Incidents/1,000 All	Incidents/1,000 Muslim > 5%	Incidents/1,000 All	Incidents/1,000 Muslim > 5%
Philippines (province/year)	Unemployment	-0.131*** (0.031)	-0.148 (0.174)	-0.152*** (0.049)	-0.113 (0.132)
	Observations	624	96	624	96
	R ²	0.33	0.35	0.141	0.354
	Controls	Ethnicity	Ethnicity	Province FE	Province FE
	DV	AGE/1,000	AGE/1,000	JOIIS/1,000	CIDNE/1,000
Afghanistan (district/month)	Unemployment	0.032** (0.018)	-0.049 (0.035)	-0.063* (0.044)	-0.049 (0.035)
	Observations	2,160	2,160	2,160	2,160
	R ²	0.17	0.08	0.105	0.075
	Controls	Ethnicity	District FE	District FE	District FE
	DV	AGE/1,000	AGE/1,000	JOIIS/1,000	CIDNE/1,000
... of which Afghanistan majority Pashtun	Unemployment	0.092** (0.054)	-0.158** (0.085)	-0.191* (0.106)	-0.163* (0.088)
	Observations	923	917	917	917
	R ²	0.10	0.22	0.25	0.21
	Controls	Ethnicity	District FE	District FE	District FE

Note: All regressions include time/wave fixed effects and are weighted by population. The Afghanistan sample has five waves of ANQAR data and all waves span two months, except wave 5 spanning one month (results are similar using wave-month fixed effects). Robust standard errors clustered by district/province reported in parentheses. Variables described in note to Table 1. Iraq has 104 districts, the Philippines have 78 provinces, and Afghanistan has 363 districts included in the samples above. *** $p < .01$. ** $p < .05$. * $p < .1$, all for one-tailed test with $H_0: \beta > 0$.

The estimated coefficients on unemployment in table 2 are almost always negative and most of them are *statistically significant*.

Recall that the opportunity cost theory predicts a positive relationship between unemployment and violence so this theory is rejected.

Now let's think a bit about the meaning of "statistically significant".

Read the footnote to the table and you see that one star corresponds to “ $p < .1$ ”, two stars to “ $p < .05$ ” and three stars to “ $p < .01$ ”.

This begs the questions: [what do the p's mean?](#)

When $p = 0.05$ this means that if we assume that the true value of the coefficient is 0 (the “null hypothesis”) then the probability that our data will lead us to estimate the (negative) coefficient that we have actually estimated or a smaller one (i.e., a more negative one) is equal to 0.05.

In other words, a low p value indicates that it is unlikely we will observe the negative value we have estimated or still lower ones if the underlying reality is that there is no relationship between unemployment and violence.

By tradition most published regression tables have the same starring system that Table 2 on slide 19 has.

More stars mean “more significant results” in the sense that we would be very unlikely to see the reported results if the null hypothesis is true, i.e., if the coefficient is really equal to 0.

The next important point to understand is that we should not stop after only assessing the statistical significance of an estimated coefficient.

We should also ask whether the coefficient has [practical significance](#).

In other words, we should also consider the magnitudes of the coefficients and decide whether they are big enough to matter. (I think that Berman et al. 1 fails to do this.)

Let's do a magnitude assessment for the very first coefficient in the table which is -1.118. This is the coefficient on unemployment in Iraq.

Here is a quick comment before proceeding. There is no standard procedure for assessing the magnitudes of regression coefficients. Doing a good job on this requires having good knowledge of the situation you are analysing and applying some common sense.

That said, I can still give you set approach that is usually helpful:

Calculate the effect of a 1 standard deviation change in a right-hand-side variable on the left-hand-side variable where the change in the left-hand-side variable is measured in units of its standard deviation. This formulation is a bit complicated so let's look now at an example.

Suppose that unemployment in a particular district for a particular quarter increases by six percentage points, say from 11.7% to 17.7%. This change is almost exactly 1 standard deviation. (Look at Table 1. Notice that an unemployment rate of, for example, 11.7 is rendered as 0.117.)

The regression coefficient we are analysing suggests that this increase in unemployment should be associated, on average, with a decrease (remember the coefficient is negative) of roughly 0.067 violent incidents per 1,000 people per quarter ($0.06 \times 1.118 = 0.067$).

Table 1 shows that the standard deviation for violent incidents per 1,000 per quarter is 0.504. 0.067 is about 13% of 0.504. *We can conclude that a 1 standard deviation change in unemployment is associated with a change of only about 13% of 1 standard deviation in the violence rate.* From this point of view the relationship between unemployment and violence seems fairly weak even though it is statistically significant.

What is the point of expressing the changes in unemployment and violence in standard deviation units?

First, a change of 1 standard deviation in the right-hand-side variable is something that can really happen, unlike a change of, say, 4 standard deviations. At the same time a 1 standard deviation change is not trivially small compared to the way this variable tends to move around. In other words, it is a moderately large change compared to the natural variation in the variable.

Similarly, if the left-hand-side variable changes by 1 standard deviation we can say that this is a moderately large change compared to the natural variation in the variable.

Nevertheless, we must bear in mind that the above technique is just one way to evaluate regression coefficient magnitudes. Some other methods can also be appropriate in some circumstances.

For example, Table 1 also says that the mean number of violent incidents per district/quarter is 0.258. 0.067 is around 26% of 0.258 ($100 \times 0.067/0.258 = 26$). So a change of 6 percentage points in the unemployment rate is associated with a change in violence that of roughly 26% of the mean level of violence, on average. Perhaps this formulation makes the effect size seem a little larger than it did under the other formulation.

In summary, I would say that the relationship between unemployment and violence in Iraq is not tiny but it is not big either.

Let's now move on to [Berman et al. 2](#) paper.

We have some familiarity with one of the Iraq data sources Berman et al. 2 use – the SIGACTS database. They use it to measure the intensity of insurgent attacks by district and time, measured now in half-year periods.

They have another key military database as well, called the “Commanders Emergency Response Program (CERP)”. These were small bits of funding available at the “brigade level or below”. (Brigades normally have something like 3,000 to 5,000 soldiers). The idea is that commanders at these relatively low levels have pots of money they can use to win friends, influence people and generally to stabilize the local situations they are involved in.

The question is – Do these strategic expenditures of funds actually help to reduce violence?

It is not too hard to imagine how one could use these datasets to address the question. In particular, we can ask whether there seems to be a negative association between insurgent attacks and CERP money spent.

However, if we want to make *causal* claims, e.g., that spending money causes violence to go down, then we have to think carefully. In particular, there are various channels of potential causation that can mess up our analysis.

Here are two plausible scenarios:

1. Insurgent attacks break out in an area and commanders quickly rush in with CERP money (and other things such as troops) to try to stabilize an area. By this description violence causes CERP spending. Therefore, we will observe a positive correlation between the two and may erroneously conclude that CERP spending causes violence.
2. Commanders do not like to spend money in unstable environments so they only spend CERP money in places that are already pretty stable. In this case lack of violence causes CERP spending and we will observe a negative correlation between violence and spending. If this scenario is important then we may conclude erroneously that CERP spending causes violence to go down

There can also be more complicated, mixed scenarios.

Variant on scenario 1. Violence attracts CERP money which, in turn, reduces violence. In this case CERP expenditure causes lower violence but this causal effect is underestimated because the CERP money only flows into environments that are violent in the first place. In fact, under this variant we may estimate either a negative or a positive correlation between CERP spending and violence, depending on how strong these two countervailing effects are.

Variant on scenario 2. CERP spending causes lower violence but we overestimate the extent to which this is true because CERP spending is only funnelled into low-violence environments in the first place.

Berman et al. 2 are not hugely concerned about the first issue on slide 30 because this problem only makes it hard for them to find that CERP spending reduces violence.

They do find a statistically significant effect of CERP spending in reducing violence and are happy to argue that maybe they have underestimated the true effect to the extent that point 1 has bite.

In other words, they are much more interested in demonstrating a statistically significant effect than they are in pinpointing its exact magnitude and the first issue on slide 30 does not throw a finding of statistical significance into question.

Point 2 is more of a challenge for Berman et al. 2 because it suggests that their findings may be overestimates. Critics could argue that if they could eliminate the source of upward bias highlighted in point 2 then their coefficient estimate may no longer be statistically significant. In fact, almost all of the authors' effort in the paper is directed at overcoming this critique.

The strategy of Berman et al. 2 is to look at *changes* in spending and insurgent attacks rather than *levels* of spending and insurgent attacks. So, although there may be many areas where spending is high and violence is low, the authors do not use this positive correlation to support a theory that CERP spending is *causing* the violence to be low. Rather, Berman et al. 2 try to make their case by showing that *increases* in CERP spending are associated with *decreases* in insurgent attacks.

In technical terms they have a [first-differences design](#). This is a sensible approach to address the causation issue however, as one student pointed out on the blog, it still leaves an issue on another level. If military commanders *increase* CERP spending to areas where violence is *decreasing*, i.e., they create positive first-differences in CERP spending where there are negative first differences in violence, then we can *erroneously* conclude that *increases* in CERP spending *cause decreases* in insurgent attacks.

Let's pause for a moment to gather some intuition.

Suppose I try to argue that having more books in a household will improve the school performance of the children living in these households. There is surely a positive correlation between school performance and books in the household but this does not mean that books *cause* high marks. For example, households with lots of books will tend to be richer than households with few books and money will come with lots of advantages for children. So the positive correlation between books and marks is not very impressive on its own.

But suppose we conduct the following experiment. Choose a bunch of household and record the school performance of all their children. Then pick some of these households at random, give them a bunch of books and record everyone's marks again. Do the marks of the children who got the books improve relative to the marks of the others? If so then it is reasonable to say that *increases* in books lead to *improvements* in marks. Berman et al. 2 are applying a very similar idea to Iraq.

Here is the equation they work with:

$$v_{i,t} - v_{i,t-1} = \beta_1(g_{i,t} - g_{i,t-1}) + \beta_2(m_{i,t} - m_{i,t-1}) + \beta_3(v_{i,t-1} - v_{i,t-2}) + \delta_t + \varepsilon_{i,t}$$

i 's are districts

t 's are half-year time periods

v 's are insurgent attacks

g 's are various measures of aid spending

m 's are measures of troop strength

δ_t is a year dummy

$\varepsilon_{i,t}$ is a random shock

So far I have written only about CERP aid because this is what Berman et. al. 2 are most interested in.

However, they also consider a variety of other aid programmes including things called the “Community Action Program (CAP)” and the “Community Stabilization Program (CSP)”. These were both administered by USAID, the main US aid agency.

I also did not mention the troop strength data above because I wanted to focus attention on the main data for the paper – SIGACTS and CERP.

Troop strength enters as a [control](#).

Why control for troop strength?

It could be that a commander increases both troop strength and CERP spending in a particular district at a particular point in time – insurgent attacks may decrease but maybe to the extent there is cause and effect here it is more the troop increase causing the decrease in violence rather than the expenditure of funds driving down the violence. By controlling for troops you can see whether there might be a spending effect above and beyond the troop effect.

Here is the main table of the paper:

TABLE 1—DEVELOPMENT PROGRAMS AND VIOLENCE SUPPRESSION: ALTERNATIVE PROGRAMS

LHS variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	Variable mean
Incidents/capita								
CERP < \$50K	-0.0639*** (0.0186)							\$1.32
CERP > \$50K		-0.0108* (0.00559)						\$8.75
Non-CERP > \$100K			0.000899 (0.000565)					\$33.18
Non-CERP < \$100K				0.00636 (0.0228)				\$0.31
CSP					-0.0470* (0.0242)			\$0.44
CAP						-0.0118 (0.0279)		\$0.18
USAID							-0.00248 (0.00334)	\$12.51
Troop strength	0.0479 (0.0349)	0.0451 (0.0400)	0.0386 (0.0325)	0.0349 (0.0337)	0.0466 (0.0328)	0.0354 (0.0333)	0.0301 (0.0339)	1.12
Lagged troop strength	0.0162 (0.0602)	-0.0373 (0.0704)	-0.0411 (0.0729)	-0.0421 (0.0732)	-0.0205 (0.0705)	-0.0424 (0.0728)	-0.0400 (0.0731)	1.13
R^2	0.222	0.204	0.180	0.179	0.194	0.179	0.180	

Notes: An observation is a district ($N=103$) \times half year. District Karkh is excluded as national CSP programs are confounded with local programs there. Means are for levels ($NT=927$) though regressions are estimated in first differences ($NT=824$). Incidents are measured per 1000 population. Their mean is 0.587. Troop strength is measured in battalions per district. Regressions are weighted by population and include year effects, Sunni vote-year interactions, lagged incidents, and a constant. Standard errors are clustered by district.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Think back to last lecture.

Standard errors are in parentheses below the coefficients. If the absolute value of the coefficient is more than two times as large as the standard error then you should expect a p value below 0.05 and you will see two stars in the table. If the absolute value of the coefficient is more than around 2.5 times the standard error then p will be below 0.01 and you will see three stars in the table.

The main take home point from the table is meant to be that small-scale CERP expenditures seem to have the biggest impact on reducing insurgent attacks. That is, by far the biggest and most significant coefficient is the -0.069 coefficient on CERP projects below \$50,000. All the other aid coefficients are either insignificant or small and of borderline significance.

But the above paragraph only looks at statistical significance. What about practical significance?

This turns out to be harder than you might expect and I had to correspond with the authors of the paper to sort it out. The table is not as clear as it could be to facilitate easy interpretation.

This experience leads me to think that many people have read the paper without really coming to grips with whether or not the estimated CERP coefficient is big enough to have practical significance. So, apparently, many people are falling into the trap of looking at statistical but not practical significance.

Let's have a look for ourselves.

The starting point is always to look at the right-hand-side variables and figure out what kinds of movements in these variables are reasonable or typical. The far-right column in this table is a good place to start because it gives means for all the variables. But the mean of CERP spending is \$1.32, a ridiculously low number, so I emailed the authors.

It turns out that this variable is spending *per capita*. They added that median population of a district is 165,000. So average (in some sense) spending was a bit over \$218,000 for an average district-half year. This is still a fairly small amount of money but not completely miniscule (like \$1.32 would be).

The authors also say that a typical small CERP project spends about \$25,000. Let's suppose that *spending in some district-period goes up by two of these projects*. That would be a fairly big change compared to an average around \$218,000 - an increase of around \$50,000 is more than 20% of \$218,000. Multiply this by -0.069 and we get *that violent incidents are reduced by 0.021*.

Is 0.021 incidents a little or a lot?

This number is, perhaps, bigger than it might seem to you at first glance since its units are incidents per capita per district-half year. So let's compare it to the mean number of incidents per district-half year per 1,000 which is 0.587. 0.021 is just 3.6% of 0.587 so the reduction really is small, just as it seemed to be at first glance.

Again, we wind up with a statistically significant effect that does not have enormous practical significance.