

Mobile safety cameras: Estimating casualty reductions

Lee Fawcett and Neil Thorpe

Karlsruhe — December 2013

Structure of this talk

- Overview/Background
- Regression to the mean
- Empirical Bayes
- Full Bayes
- Summary

Structure of this talk

- Overview/Background
- Regression to the mean
- Empirical Bayes
- Full Bayes
- Summary

Structure of this talk

- Overview/Background
- Regression to the mean
- Empirical Bayes
- Full Bayes
- Summary

Structure of this talk

- Overview/Background
- Regression to the mean
- Empirical Bayes
- Full Bayes
- Summary

Structure of this talk

- Overview/Background
- Regression to the mean
- Empirical Bayes
- Full Bayes
- Summary

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000 paper:** Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000 paper:** Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000 paper:** Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000** paper: Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000** paper: Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

- **1996:** Government report: road safety cameras effective weapon in reducing casualty figures
- High implementation/running costs
- **1998:** Government allowed traffic authorities to recover these costs via speeding fines
- **2000** paper: Speed cameras an important part of the government's 2010 casualty reduction targets
- **April 2000:** two year pilot programme involving eight road safety camera partnerships (SCPs)
- Results at the end of 2000 prompted an earlier-than-expected national roll-out of SCPs

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

NSCP

- Joined the national programme in April 2003
- 56 mobile speed camera sites – form our **treated set**
- 'before' period (April 2001–March 2003) vs 'after' period (April 2004–March 2006)

Aims:

- To investigate changes in the number/severity of casualties
- To investigate changes in cost-of-treatment estimates

Speed cameras: Friend or foe?



Speed cameras: Friend or foe?



Speed cameras: Friend or foe?



Speed cameras: Friend or foe?



Regression To *[the]* Mean (RTM)

The 56 sites were chosen because of their unusually high casualty history (“blackspots”).

- The number of casualties is bound to decrease in any ‘after’ period, “...even if a garden gnome is used instead of a speed camera” (*Paul Smith, SafeSpeed*)



- Main consequence: ‘before’ versus ‘after’ will probably exaggerate the treatment effect
- Studies have shown that a reduction owing to RTM of between 20–30% is common

Regression To *[the]* Mean (RTM)

The 56 sites were chosen because of their unusually high casualty history (“blackspots”).

- The number of casualties is bound to decrease in any ‘after’ period, “...even if a garden gnome is used instead of a speed camera” (*Paul Smith, SafeSpeed*)



- Main consequence: ‘before’ versus ‘after’ will probably exaggerate the treatment effect
- Studies have shown that a reduction owing to RTM of between 20–30% is common

Regression To *[the]* Mean (RTM)

The 56 sites were chosen because of their unusually high casualty history (“blackspots”).

- The number of casualties is bound to decrease in any ‘after’ period, “...even if a garden gnome is used instead of a speed camera” (*Paul Smith, SafeSpeed*)



- Main consequence: ‘before’ versus ‘after’ will probably exaggerate the treatment effect
- Studies have shown that a reduction owing to RTM of between 20–30% is common

Regression To *[the]* Mean (RTM)

The 56 sites were chosen because of their unusually high casualty history (“blackspots”).

- The number of casualties is bound to decrease in any ‘after’ period, “...even if a garden gnome is used instead of a speed camera” (*Paul Smith, SafeSpeed*)



- Main consequence: ‘before’ versus ‘after’ will probably exaggerate the treatment effect
- Studies have shown that a reduction owing to RTM of between 20–30% is common

Regression To *[the]* Mean (RTM)


The 56 sites were chosen because of their unusually high casualty history (“blackspots”).

- The number of casualties is bound to decrease in any ‘after’ period, “...even if a garden gnome is used instead of a speed camera” (*Paul Smith, SafeSpeed*)




- Main consequence: ‘before’ versus ‘after’ will probably exaggerate the treatment effect
- Studies have shown that a reduction owing to RTM of between 20–30% is common


Regression To *[the]* Mean (RTM)




Solution: Bayesian approach



Solution: Bayesian approach



Solution: Bayesian approach



Solution: Bayesian approach

So:

- For each site, the observed casualty frequency in the before period gives us the “**likelihood**”
- We need a way of getting the “**prior**” value – this can be thought of as the casualty frequency that we’d expect to see at each site *ordinarily*
- **Bayes’ Theorem** is then used to bring the two together – giving a “**posterior**” which is a synthesis of the prior and the likelihood
- Without this, any assessment of speed camera effectiveness is likely to be exaggerated

Solution: Bayesian approach

So:

- For each site, the observed casualty frequency in the before period gives us the “**likelihood**”
- We need a way of getting the “**prior**” value – this can be thought of as the casualty frequency that we'd expect to see at each site *ordinarily*
- **Bayes' Theorem** is then used to bring the two together – giving a “**posterior**” which is a synthesis of the prior and the likelihood
- Without this, any assessment of speed camera effectiveness is likely to be exaggerated

Solution: Bayesian approach

So:

- For each site, the observed casualty frequency in the before period gives us the “**likelihood**”
- We need a way of getting the “**prior**” value – this can be thought of as the casualty frequency that we’d expect to see at each site *ordinarily*
- Bayes’ Theorem is then used to bring the two together – giving a “**posterior**” which is a synthesis of the prior and the likelihood
- Without this, any assessment of speed camera effectiveness is likely to be exaggerated

Solution: Bayesian approach

So:

- For each site, the observed casualty frequency in the before period gives us the “**likelihood**”
- We need a way of getting the “**prior**” value – this can be thought of as the casualty frequency that we’d expect to see at each site *ordinarily*
- **Bayes’ Theorem** is then used to bring the two together – giving a “**posterior**” which is a synthesis of the prior and the likelihood
- Without this, any assessment of speed camera effectiveness is likely to be exaggerated

Solution: Bayesian approach

So:

- For each site, the observed casualty frequency in the before period gives us the “**likelihood**”
- We need a way of getting the “**prior**” value – this can be thought of as the casualty frequency that we’d expect to see at each site *ordinarily*
- **Bayes’ Theorem** is then used to bring the two together – giving a “**posterior**” which is a synthesis of the prior and the likelihood
- Without this, any assessment of speed camera effectiveness is likely to be exaggerated

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach

One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Solution: Bayesian approach


One way of obtaining the prior value for each site is to look at what has happened at other sites during the before period.

- Call these other sites the **reference** set
- Need the reference sites to be **representative** of the treated sites:
 - Similar speed limits
 - Similar *observed* speeds
 - Similar mix of road type/classification
 - Similar traffic flows
 - ...
 - But **not** similar casualty frequencies! Want sites with more “normal” casualty frequencies!
- Build a **casualty prediction model** using data from the reference set, and then apply it to the treated set

Analysis I

The Empirical Bayes approach

Step 1: Do we even have a RTM problem?



Step 1: Do we even have a RTM problem?

More formal test: **Kendall's turning point test.**

- Hypothesis H_0 : Casualty frequencies are unsystematic (i.e. this *is* just a 'blip'!)
- Compares the number of "turning points" to what we'd expect to see under random variation
- For our data: **Retain** H_0 ...
- ... looks like we are at the peak of a 'blip'!

Step 1: Do we even have a RTM problem?

More formal test: **Kendall's turning point test.**

- Hypothesis H_0 : Casualty frequencies are unsystematic (i.e. this *is* just a 'blip'!)
- Compares the number of "turning points" to what we'd expect to see under random variation
- For our data: **Retain H_0** ...
- ... looks like we are at the peak of a 'blip'!

Step 1: Do we even have a RTM problem?

More formal test: **Kendall's turning point test.**

- Hypothesis H_0 : Casualty frequencies are unsystematic (i.e. this *is* just a 'blip'!)
- Compares the number of "turning points" to what we'd expect to see under random variation
- For our data: **Retain H_0** ...
- ... looks like we are at the peak of a 'blip'!

Step 1: Do we even have a RTM problem?

More formal test: **Kendall's turning point test.**

- Hypothesis H_0 : Casualty frequencies are unsystematic (i.e. this *is* just a 'blip'!)
- Compares the number of "turning points" to what we'd expect to see under random variation
- For our data: **Retain H_0** ...
- ... looks like we are at the peak of a 'blip'!

Step 1: Do we even have a RTM problem?

More formal test: **Kendall's turning point test**.

- Hypothesis H_0 : Casualty frequencies are unsystematic (i.e. this *is* just a 'blip'!)
- Compares the number of "turning points" to what we'd expect to see under random variation
- For our data: **Retain** H_0 ...
- ... looks like we are at the peak of a 'blip'!

Step 2: Is the reference set representative?

From 67 (non-speed camera) sites in Northumbria, we have data relating to:

- x_1 : Speed limit (mph)
- x_2 : Average observed speed (mph)
- x_3 : 85th percentile speed (mph)
- x_4 : % of drivers over the speed limit
- x_5 : % of drivers at least 15mph over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)
- x_8 : Road type (single/dual/mixed)

Step 2: Is the reference set representative?

From 67 (non-speed camera) sites in Northumbria, we have data relating to:

- x_1 : Speed limit (mph)
- x_2 : Average observed speed (mph)
- x_3 : 85th percentile speed (mph)
- x_4 : % of drivers over the speed limit
- x_5 : % of drivers at least 15mph over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)
- x_8 : Road type (single/dual/mixed)

Step 2: Is the reference set representative?

Principal Components Analysis/Multiple Factor Analysis on

		X_1	X_2	...	variable p	...	X_8	
Reference	site 1							}
	⋮	⋮	⋮	...	...	...	⋮	
	site i	⋮	⋮	...	...	...	⋮	
	⋮	⋮	⋮	...	...	...	...	
	site 67							
Treated	site 1							
	⋮							
	site j	⋮	⋮	...	...	...	⋮	
	⋮	⋮	⋮	...	...	...	⋮	
	site 56							

Plot scores for first PC against those for second PC, using different plotting character for "reference" and "treated".

Step 2: Is the reference set representative?

Principal Components Analysis/Multiple Factor Analysis on

		X_1	X_2	$\dots$	variable p	$\dots$	X_8	
Reference	site 1							} X
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	site i	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\dots$	
	site 67							
Treated	site 1							
	$\vdots$							
	site j	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	site 56							

Plot scores for first PC against those for second PC, using different plotting character for "reference" and "treated".

Step 2: Is the reference set representative?


Principal Components Analysis/Multiple Factor Analysis on

		X_1	X_2	$\dots$	variable p	$\dots$	X_8	
Reference	site 1							}
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	site i	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\dots$	
	site 67							
Treated	site 1							
	$\vdots$							
	site j	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	$\vdots$	$\vdots$	$\vdots$	$\dots$	$\dots$	$\dots$	$\vdots$	
	site 56							

Plot scores for first PC against those for second PC, using different plotting character for “reference” and “treated”.

Step 2: Is the reference set representative?

Multiple factor analysis on X



Step 2: Is the reference set representative?

More formally – **Permutation tests**:

- The **Mahalanobis distance** can be thought of as a measure of “mixing”
- Compare the Mahalanobis distance to that obtained under a random allocation of sites to both treated and reference sets
- If the points “mix well”, then there will be no significant difference between the *real* Mahalanobis distance and those obtained after randomisation
- For our data: points mix well!

Step 2: Is the reference set representative?

More formally – **Permutation tests**:

- The **Mahalanobis distance** can be thought of as a measure of “mixing”
- Compare the Mahalanobis distance to that obtained under a random allocation of sites to both treated and reference sets
- If the points “mix well”, then there will be no significant difference between the *real* Mahalanobis distance and those obtained after randomisation
- For our data: points mix well!

Step 2: Is the reference set representative?

More formally – **Permutation tests**:

- The **Mahalanobis distance** can be thought of as a measure of “mixing”
- Compare the Mahalanobis distance to that obtained under a random allocation of sites to both treated and reference sets
- If the points “mix well”, then there will be no significant difference between the *real* Mahalanobis distance and those obtained after randomisation
- For our data: points mix well!

Step 2: Is the reference set representative?

More formally – **Permutation tests**:

- The **Mahalanobis distance** can be thought of as a measure of “mixing”
- Compare the Mahalanobis distance to that obtained under a random allocation of sites to both treated and reference sets
- If the points “mix well”, then there will be no significant difference between the *real* Mahalanobis distance and those obtained after randomisation
- For our data: points mix well!

Step 2: Is the reference set representative?

More formally – **Permutation tests**:

- The **Mahalanobis distance** can be thought of as a measure of “mixing”
- Compare the Mahalanobis distance to that obtained under a random allocation of sites to both treated and reference sets
- If the points “mix well”, then there will be no significant difference between the *real* Mahalanobis distance and those obtained after randomisation
- For our data: points mix well!

Step 2: Is the reference set representative?

Even more **Permutation tests** – compare individual variables between treatment and reference sites using:

$$\delta_p = |\bar{x}_p^{\text{TRT}} - \bar{x}_p^{\text{REF}}|, \quad p = 1, \dots, 8.$$

If the treatment and reference sites *are* representative, then the values of δ_p would not be significantly different to those obtained after a random allocation of sites to each set.

Step 2: Is the reference set representative?

Even more **Permutation tests** – compare individual variables between treatment and reference sites using:

$$\delta_p = |\bar{x}_p^{\text{TRT}} - \bar{x}_p^{\text{REF}}|, \quad p = 1, \dots, 8.$$

If the treatment and reference sites *are* representative, then the values of δ_p would not be significantly different to those obtained after a random allocation of sites to each set.

Step 2: Is the reference set representative?

Even more **Permutation tests** – compare individual variables between treatment and reference sites using:

$$\delta_p = |\bar{x}_p^{\text{TRT}} - \bar{x}_p^{\text{REF}}|, \quad p = 1, \dots, 8.$$

If the treatment and reference sites *are* representative, then the values of δ_p would not be significantly different to those obtained after a random allocation of sites to each set.

Step 2: Is the reference set representative?

- Hypothesis H_0 : Variables are exchangeable
- Randomly allocate sites to each set, lots of times! Each time, find δ_p , $p = 1, \dots, 8$
- Compare the *real* δ_p to those obtained under random allocation
- For our data: **Retain** H_0 for all 8 variables
- If H_0 is rejected, for which variable? Need to reconsider the reference sites – especially for those variables which “failed”

Step 2: Is the reference set representative?

- Hypothesis H_0 : Variables are exchangeable
- Randomly allocate sites to each set, lots of times! Each time, find $\delta_p, p = 1, \dots, 8$
- Compare the *real* δ_p to those obtained under random allocation
- For our data: **Retain** H_0 for all 8 variables
- If H_0 is rejected, for which variable? Need to reconsider the reference sites – especially for those variables which “failed”

Step 2: Is the reference set representative?

- Hypothesis H_0 : Variables are exchangeable
- Randomly allocate sites to each set, lots of times! Each time, find $\delta_p, p = 1, \dots, 8$
- Compare the *real* δ_p to those obtained under random allocation
- For our data: **Retain** H_0 for all 8 variables
- If H_0 is rejected, for which variable? Need to reconsider the reference sites – especially for those variables which “failed”

Step 2: Is the reference set representative?

- Hypothesis H_0 : Variables are exchangeable
- Randomly allocate sites to each set, lots of times! Each time, find $\delta_p, p = 1, \dots, 8$
- Compare the *real* δ_p to those obtained under random allocation
- For our data: **Retain** H_0 for all 8 variables
- If H_0 is rejected, for which variable? Need to reconsider the reference sites – especially for those variables which “failed”

Step 2: Is the reference set representative?

- Hypothesis H_0 : Variables are exchangeable
- Randomly allocate sites to each set, lots of times! Each time, find $\delta_p, p = 1, \dots, 8$
- Compare the *real* δ_p to those obtained under random allocation
- For our data: **Retain** H_0 for all 8 variables
- If H_0 is rejected, for which variable? Need to reconsider the reference sites – especially for those variables which “failed”

Step 3: Constructing the prior

From 67 (non-speed camera) sites in Northumbria, we were given data relating to:

- x_1 : Speed limit (mph)
- x_2 : Average observed speed (mph)
- x_3 : 85th percentile speed (mph)
- x_4 : % of drivers over the speed limit
- x_5 : % of drivers at least 15mph over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)
- x_8 : Road type (single/dual/mixed)

Standard regression techniques were used to obtain the casualty prediction model.

Step 3: Constructing the prior

From 67 (non-speed camera) sites in Northumbria, we were given data relating to:

- x_1 : Speed limit (mph)
- x_2 : Average observed speed (mph)
- x_3 : 85th percentile speed (mph)
- x_4 : % of drivers over the speed limit
- x_5 : % of drivers at least 15mph over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)
- x_8 : Road type (single/dual/mixed)

Standard regression techniques were used to obtain the casualty prediction model.

Step 3: Constructing the prior

From 67 (non-speed camera) sites in Northumbria, we were given data relating to:

- x_1 : Speed limit (mph)
- x_2 : Average observed speed (mph)
- x_3 : 85th percentile speed (mph)
- x_4 : % of drivers over the speed limit
- x_5 : % of drivers at least 15mph over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)
- x_8 : Road type (single/dual/mixed)

Standard regression techniques were used to obtain the casualty prediction model.

Step 3: Constructing the prior

- x_2 : Average observed speed (mph)
- x_4 : % of drivers over the speed limit
- x_6 : Daily traffic flow
- x_7 : Road classification (A, B, C, U)

Step 3: Constructing the prior

This gives:

$$\hat{\mu}_i = \exp \{ 1.93 - 0.04x_{2_i} - 0.01x_{4_i} + 0.44x_{6_i} + 0.67I_{1_i} + 0.85I_{2_i} + 1.06I_{3_i} \}$$

which is then applied to variables x_2 , x_4 , x_6 and x_7 from the

Step 3: Constructing the prior

This gives:

$$\hat{\mu}_j = \exp \{ 1.93 - 0.04x_{2_j} - 0.01x_{4_j} + 0.44x_{6_j} + 0.67l_{1_j} + 0.85l_{2_j} + 1.06l_{3_j} \}$$

which is then applied to variables x_2 , x_4 , x_7 and x_8 from the **treated set**.

Step 4: Obtaining the posterior

The prior value (μ_j) is combined with the actual observed value ($y_{j,\text{before}}$) to get the **posterior** estimate of casualty frequency at each site.

In the **Empirical Bayes** (EB) approach, this is given as:

$$EB_j = \alpha_j \mu_j + (1 - \alpha_j) y_{j,\text{before}},$$

that is, a weighted sum of what we actually observed, and what we'd expect to see normally, at each treated site j .

Step 4: Obtaining the posterior

The prior value (μ_j) is combined with the actual observed value ($y_{j,\text{before}}$) to get the **posterior** estimate of casualty frequency at each site.

In the **Empirical Bayes** (EB) approach, this is given as:

$$EB_j = \alpha_j \mu_j + (1 - \alpha_j) y_{j,\text{before}},$$

that is, a weighted sum of what we actually observed, and what we'd expect to see normally, at each treated site j .

Step 4: Obtaining the posterior

The prior value (μ_j) is combined with the actual observed value ($y_{j,\text{before}}$) to get the **posterior** estimate of casualty frequency at each site.

In the **Empirical Bayes** (EB) approach, this is given as:

$$EB_j = \alpha_j \mu_j + (1 - \alpha_j) y_{j,\text{before}},$$

that is, a weighted sum of what we actually observed, and what we'd expect to see normally, at each treated site j .

Step 4: Obtaining the posterior

The prior value (μ_j) is combined with the actual observed value ($y_{j,\text{before}}$) to get the **posterior** estimate of casualty frequency at each site.


In the **Empirical Bayes** (EB) approach, this is given as:

$$EB_j = \alpha_j \mu_j + (1 - \alpha_j) y_{j,\text{before}},$$

that is, a weighted sum of what we actually observed, and what we'd expect to see normally, at each treated site j .

Step 4: Obtaining the posterior

Mathematically, this is the mean of the posterior in the picture we saw before:



Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
⋮	⋮	⋮	⋮	⋮	⋮
Site 13	12	1.71	0.59	5.95	2
⋮	⋮	⋮	⋮	⋮	⋮
Site 47	16	7.84	0.24	14.06	5
⋮	⋮	⋮	⋮	⋮	⋮
Total	436			321	298

Site 13: Observed change: -10; after RTM: -4

Total: Observed change: -138; after RTM: -23 → 26.4% RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10 ; after RTM: -4

Total: Observed change: -138 ; after RTM: $-23 \rightarrow 26.4\%$ RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10 ; after RTM: -4

Total: Observed change: -138 ; after RTM: $-23 \rightarrow 26.4\%$ RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10 ; after RTM: -4

Total: Observed change: -138 ; after RTM: $-23 \rightarrow 26.4\%$ RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10; after RTM: -4

Total: Observed change: -138; after RTM: -23 → 26.4% RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10; after RTM: -4

Total: Observed change: -138; after RTM: -23 → 26.4% RTM.

Some results

	$y_{j,\text{before}}$	μ_j	α_j	EB	$y_{j,\text{after}}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 13	12	1.71	0.59	5.95	2
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Site 47	16	7.84	0.24	14.06	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
Total	436			321	298

Site 13: Observed change: -10 ; after RTM: -4

Total: Observed change: -138 ; after RTM: $-23 \rightarrow 26.4\%$ RTM.

- We only get the “mean” of what we’d expect to see if no camera had been installed
- Perhaps some other summary would be more appropriate
- Once we’d estimated the casualty prediction model, it was treated as “perfect” – no allowances made for estimation uncertainty:

$$\hat{\mu}_j = \exp \{ 1.93 - 0.04x_{2j} - 0.01x_{4j} + 0.44x_{6j} + 0.67I_{1j} + 0.85I_{2j} + 1.06I_{3j} \}$$

- A **Fully Bayesian** (FB) analysis provides a way forward here

- We only get the “mean” of what we’d expect to see if no camera had been installed
- Perhaps some other summary would be more appropriate
- Once we’d estimated the casualty prediction model, it was treated as “perfect” – no allowances made for estimation uncertainty:

$$\hat{\mu}_j = \exp \{ 1.93 - 0.04x_{2j} - 0.01x_{4j} + 0.44x_{6j} + 0.67I_{1j} + 0.85I_{2j} + 1.06I_{3j} \}$$

- A **Fully Bayesian** (FB) analysis provides a way forward here

- We only get the “mean” of what we’d expect to see if no camera had been installed
- Perhaps some other summary would be more appropriate
- Once we’d estimated the casualty prediction model, it was treated as “perfect” – no allowances made for estimation uncertainty:

$$\hat{\mu}_j = \exp \{ 1.93 - 0.04x_{2_j} - 0.01x_{4_j} + 0.44x_{6_j} + 0.67l_{1_j} + 0.85l_{2_j} + 1.06l_{3_j} \}$$

- A **Fully Bayesian** (FB) analysis provides a way forward here

Analysis II

The Full Bayes approach

Step 1: Do we even have a RTM problem?


As before...

Step 2: Is the reference set representative?

Also, as before...

Steps 3 and 4: Prior and posterior

The aim now is to somehow obtain the full posterior distribution, not just the mean:



Steps 3 and 4: Prior and posterior

We also want to acknowledge that the casualty prediction model is an **estimate**, and will not be perfect:

$$\hat{\mu}_i = \exp \{ 1.93 - 0.04x_{2_i} - 0.01x_{4_i} + 0.44x_{6_i} + 0.67l_{1_i} + 0.85l_{2_i} + 1.06l_{3_i} \}$$

Steps 3 and 4: Prior and posterior

We also want to acknowledge that the casualty prediction model is an **estimate**, and will not be perfect:

$$\hat{\mu}_i = \exp \{ \beta_0 + \beta_1 x_{2,i} + \beta_2 x_{4,i} + \beta_3 x_{6,i} + \beta_4 l_{1,i} + \beta_5 l_{2,i} + \beta_6 l_{3,i} \}$$

We acknowledge that the coefficients are estimates by building into the analysis an accepted level of “variability” for the coefficients.

Steps 3 and 4: Prior and posterior

We also want to acknowledge that the casualty prediction model is an **estimate**, and will not be perfect:

$$\hat{\mu}_i = \exp \{ \beta_0 + \beta_1 x_{2_i} + \beta_2 x_{4_i} + \beta_3 x_{6_i} + \beta_4 l_{1_i} + \beta_5 l_{2_i} + \beta_6 l_{3_i} \}$$

We acknowledge that the coefficients are estimates by building into the analysis an accepted level of “variability” for the coefficients.

Steps 3 and 4: Prior and posterior

We also want to acknowledge that the casualty prediction model is an **estimate**, and will not be perfect:

$$\hat{\mu}_i = \exp \{ \beta_0 + \beta_1 x_{2_i} + \beta_2 x_{4_i} + \beta_3 x_{6_i} + \beta_4 l_{1_i} + \beta_5 l_{2_i} + \beta_6 l_{3_i} \}$$

We acknowledge that the coefficients are estimates by building into the analysis an accepted level of “variability” for the coefficients.

Steps 3 and 4: Prior and posterior

We also want to acknowledge that the casualty prediction model is an **estimate**, and will not be perfect:

$$\hat{\mu}_i = \exp \{ \beta_0 + \beta_1 x_{2_i} + \beta_2 x_{4_i} + \beta_3 x_{6_i} + \beta_4 l_{1_i} + \beta_5 l_{2_i} + \beta_6 l_{3_i} \}$$

We acknowledge that the coefficients are estimates by building into the analysis an accepted level of “variability” for the coefficients.

Steps 3 and 4: Prior and posterior

Advantages:

- This is the right thing to do – mathematically, but also scientifically!
- This is the aim of **any** Bayesian analysis – it's just that when these ideas were implemented in accident analysis, the computing power wasn't available to do the calculations!
- Gives a more realistic – and fair – assessment of the effectiveness of a road safety scheme
- Gives us access to the full “distribution” of what we'd expect to see if no safety scheme implemented

Steps 3 and 4: Prior and posterior

Advantages:

- This is the right thing to do – mathematically, but also scientifically!
- This is the aim of **any** Bayesian analysis – it's just that when these ideas were implemented in accident analysis, the computing power wasn't available to do the calculations!
- Gives a more realistic – and fair – assessment of the effectiveness of a road safety scheme
- Gives us access to the full “distribution” of what we'd expect to see if no safety scheme implemented

Steps 3 and 4: Prior and posterior

Advantages:

- This is the right thing to do – mathematically, but also scientifically!
- This is the aim of **any** Bayesian analysis – it's just that when these ideas were implemented in accident analysis, the computing power wasn't available to do the calculations!
- Gives a more realistic – and fair – assessment of the effectiveness of a road safety scheme
- Gives us access to the full “distribution” of what we'd expect to see if no safety scheme implemented

Steps 3 and 4: Prior and posterior

Advantages:

- This is the right thing to do – mathematically, but also scientifically!
- This is the aim of **any** Bayesian analysis – it's just that when these ideas were implemented in accident analysis, the computing power wasn't available to do the calculations!
- Gives a more realistic – and fair – assessment of the effectiveness of a road safety scheme
- Gives us access to the full “distribution” of what we'd expect to see if no safety scheme implemented

Steps 3 and 4: Prior and posterior


Advantages:

- This is the right thing to do – mathematically, but also scientifically!
- This is the aim of **any** Bayesian analysis – it's just that when these ideas were implemented in accident analysis, the computing power wasn't available to do the calculations!
- Gives a more realistic – and fair – assessment of the effectiveness of a road safety scheme
- Gives us access to the full “distribution” of what we'd expect to see if no safety scheme implemented

Some results

	Posterior			
	Mean	St. dev.	Median	95% credible interval
⋮	⋮	⋮	⋮	⋮
Site 2	2.47	1.19	2.26	(0.78, 5.37)
	<i>2.38</i>	<i>0.936</i>	—	—
⋮	⋮	⋮	⋮	⋮
Site 13	6.28	2.45	5.945	(2.50, 12.04)
	<i>5.95</i>	<i>1.56</i>	—	—
⋮	⋮	⋮	⋮	⋮
Site 47	14.23	3.47	11.03	(8.32, 21.84)
	<i>14.06</i>	<i>3.27</i>	—	—
⋮	⋮	⋮	⋮	⋮
Total T	322	23.83	308	(289.92, 369.97)

Some results



Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- "Raw" difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: **(£72,600, £4.2 million)**

Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- "Raw" difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: **(£72,600, £4.2 million)**

Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- “Raw” difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: **(£72,600, £4.2 million)**

Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- “Raw” difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: **(£72,600, £4.2 million)**


Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- “Raw” difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: (£72,600, £4.2 million)

Some results

- Total casualties before speed cameras: **436**
- Total casualties after speed cameras: **298**
- “Raw” difference: **138**
- We would expect about **308** casualties after RTM, giving a more realistic reduction of just **10** casualties
- After accounting for general trends, this reduces further to a reduction of just **8** casualties owing to the speed cameras; our uncertainty can be quantified in the FB approach
- Corresponding financial savings: **(£72,600, £4.2 million)**

Summary: General algorithm



Summary: Computing

- The statistical software package **R** was used for all the analysis:
 - Freely available
 - Flexible: can write your own functions but also has some intrinsic math/stat functions; can also write/install “packages”
 - Not user-friendly!
- All steps (1–4) easy to implement for EB; step 4 in FB is computationally expensive
- There is a clear need for a package that pulls all of this together – and makes it easy for practitioners to use!

Summary: Computing

- The statistical software package **R** was used for all the analysis:
 - Freely available
 - Flexible: can write your own functions but also has some intrinsic math/stat functions; can also write/install “packages”
 - Not user-friendly!
- All steps (1–4) easy to implement for EB; step 4 in FB is computationally expensive
- There is a clear need for a package that pulls all of this together – and makes it easy for practitioners to use!

Summary: Computing

- The statistical software package **R** was used for all the analysis:
 - Freely available
 - Flexible: can write your own functions but also has some intrinsic math/stat functions; can also write/install “packages”
 - Not user-friendly!
- All steps (1–4) easy to implement for EB; step 4 in FB is computationally expensive
- There is a clear need for a package that pulls all of this together – and makes it easy for practitioners to use!

Summary: Computing

- The statistical software package **R** was used for all the analysis:
 - Freely available
 - Flexible: can write your own functions but also has some intrinsic math/stat functions; can also write/install “packages”
 - Not user-friendly!
- All steps (1–4) easy to implement for EB; step 4 in FB is computationally expensive
- There is a clear need for a package that pulls all of this together – and makes it easy for practitioners to use!

Summary: Computing

- The statistical software package **R** was used for all the analysis:
 - Freely available
 - Flexible: can write your own functions but also has some intrinsic math/stat functions; can also write/install “packages”
 - Not user-friendly!
- All steps (1–4) easy to implement for EB; step 4 in FB is computationally expensive
- There is a clear need for a package that pulls all of this together – and makes it easy for practitioners to use!

Summary: Overview (again)

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

Summary: Overview (again)

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

Summary: Overview (again)

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

Summary: Overview (again)

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers

Summary: Overview (again)

- New procedure for evaluating the impact of road safety interventions (see Thorpe & Fawcett (2012) and Fawcett & Thorpe (2013))
- More accurate treatment of Regression To (the) Mean effect
- Developed during a case study of speed cameras in the UK
- Applicable to any site-based road safety interventions
- Incorporate into software as an evaluation tool for road safety engineers