Speed accuracy trade-off

HMM for mixed data, covariate on transition probabilities, parameter constraints

Ingmar Visser & Maarten Speekenbrink

Speed accuracy trade-off task

1. lexical decision: word or non-word
2. reward for speed versus accuracy changes on a trial-by-trial basis
3. is the trade-off discrete of gradual?

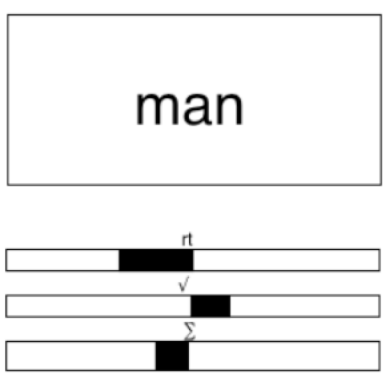

Speed accuracy trade-off data

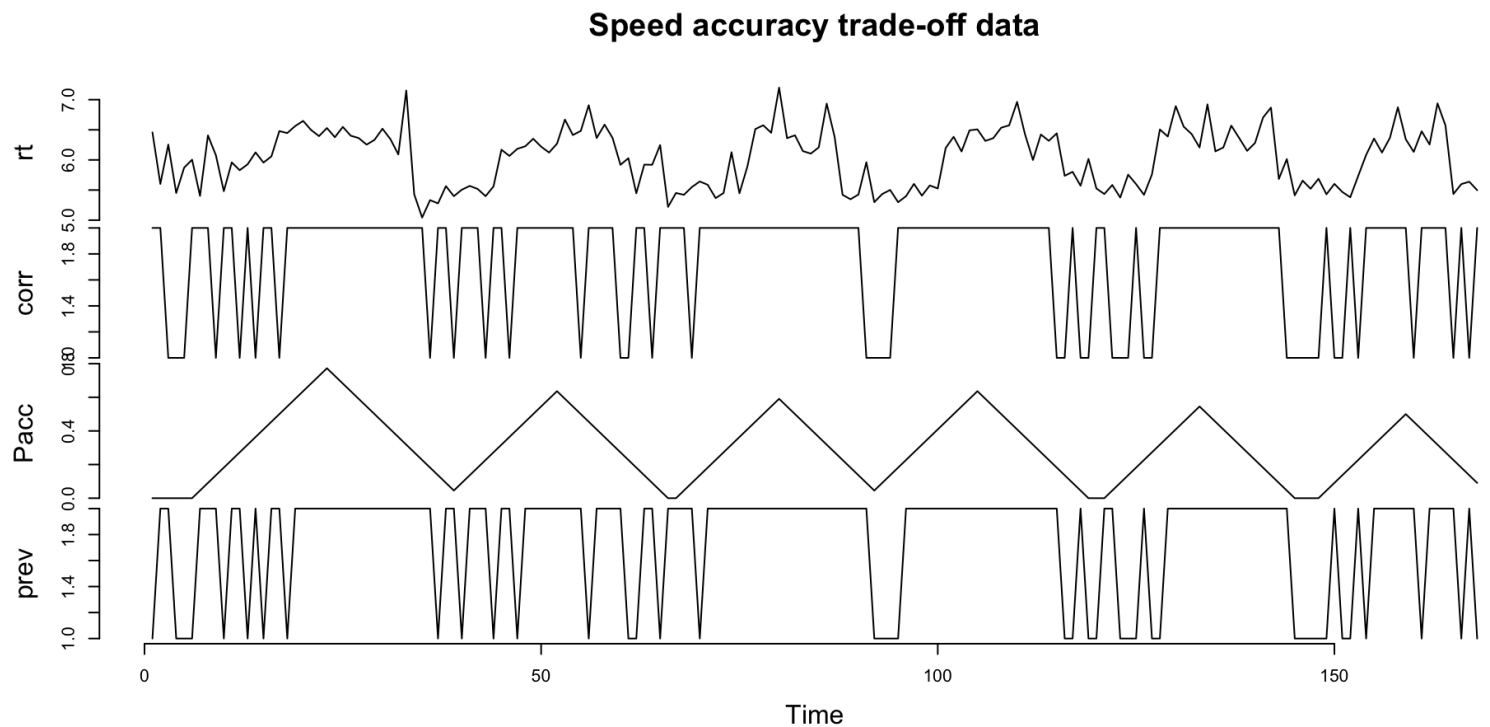

1. What is the relationship between Pacc, RT and accuracy?

2. Is the trade-off continuous or discontinuous?

Speed accuracy trade-off

Preliminaries: data bi- or multi-modal?

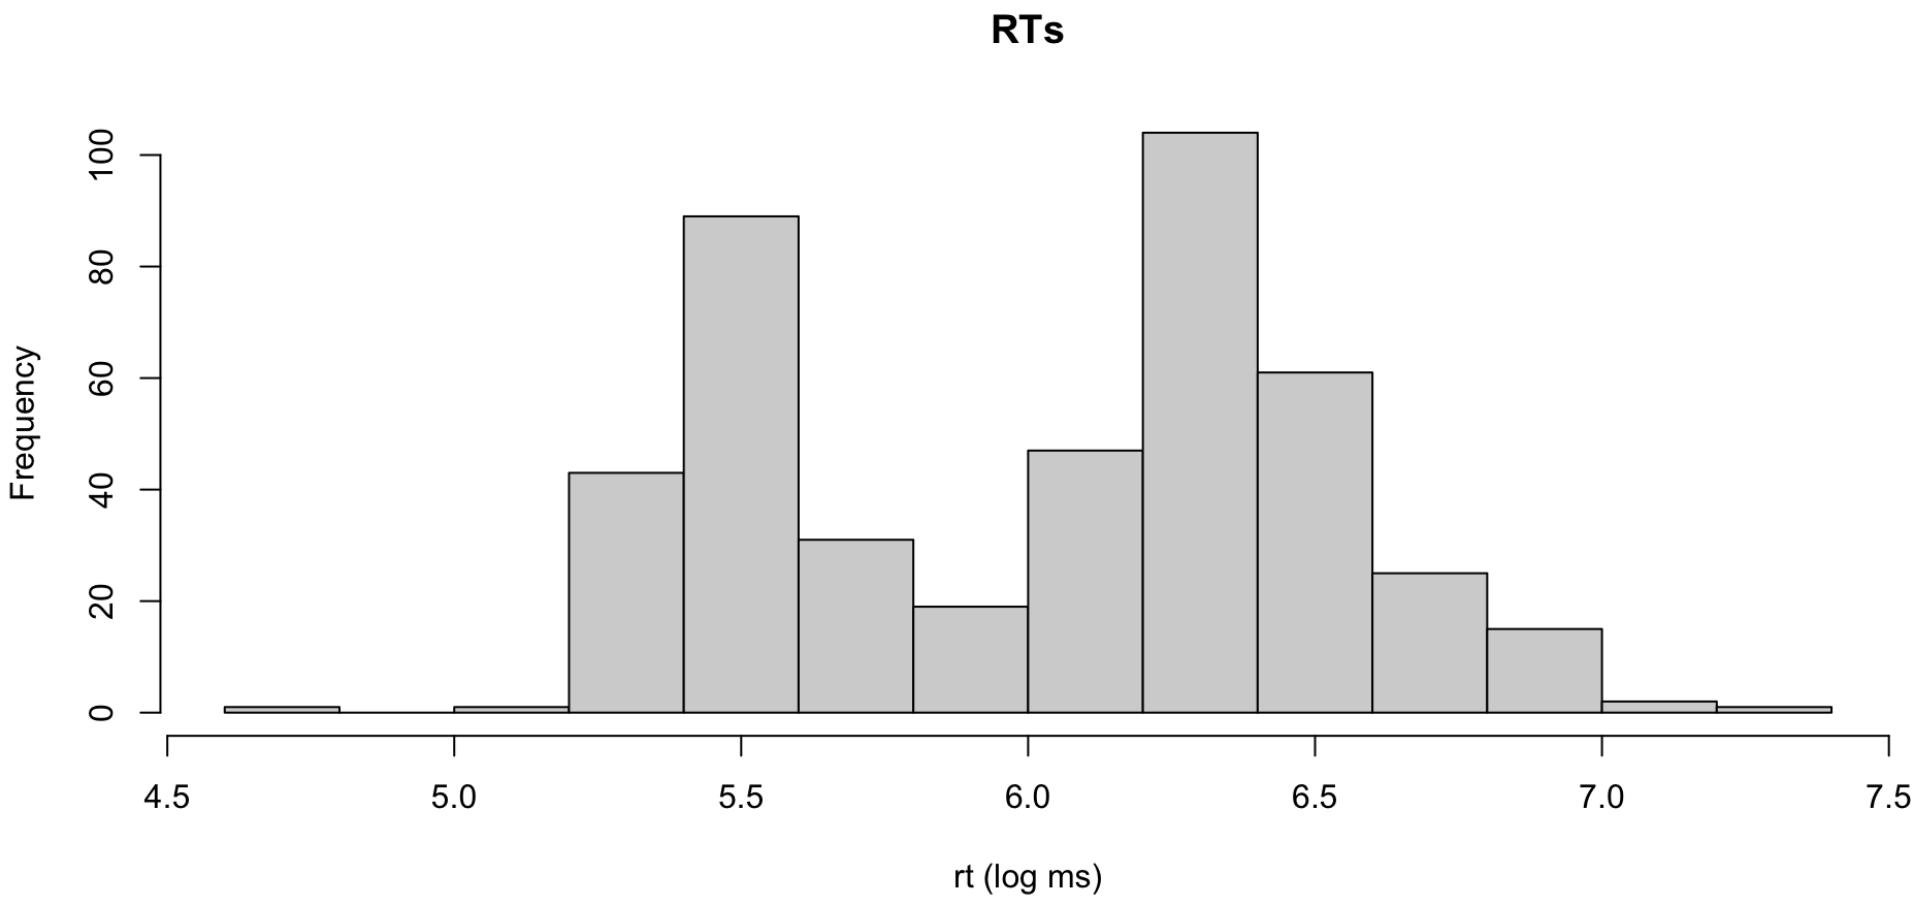

Speed accuracy trade-off

Preliminaries: dependence in the data?

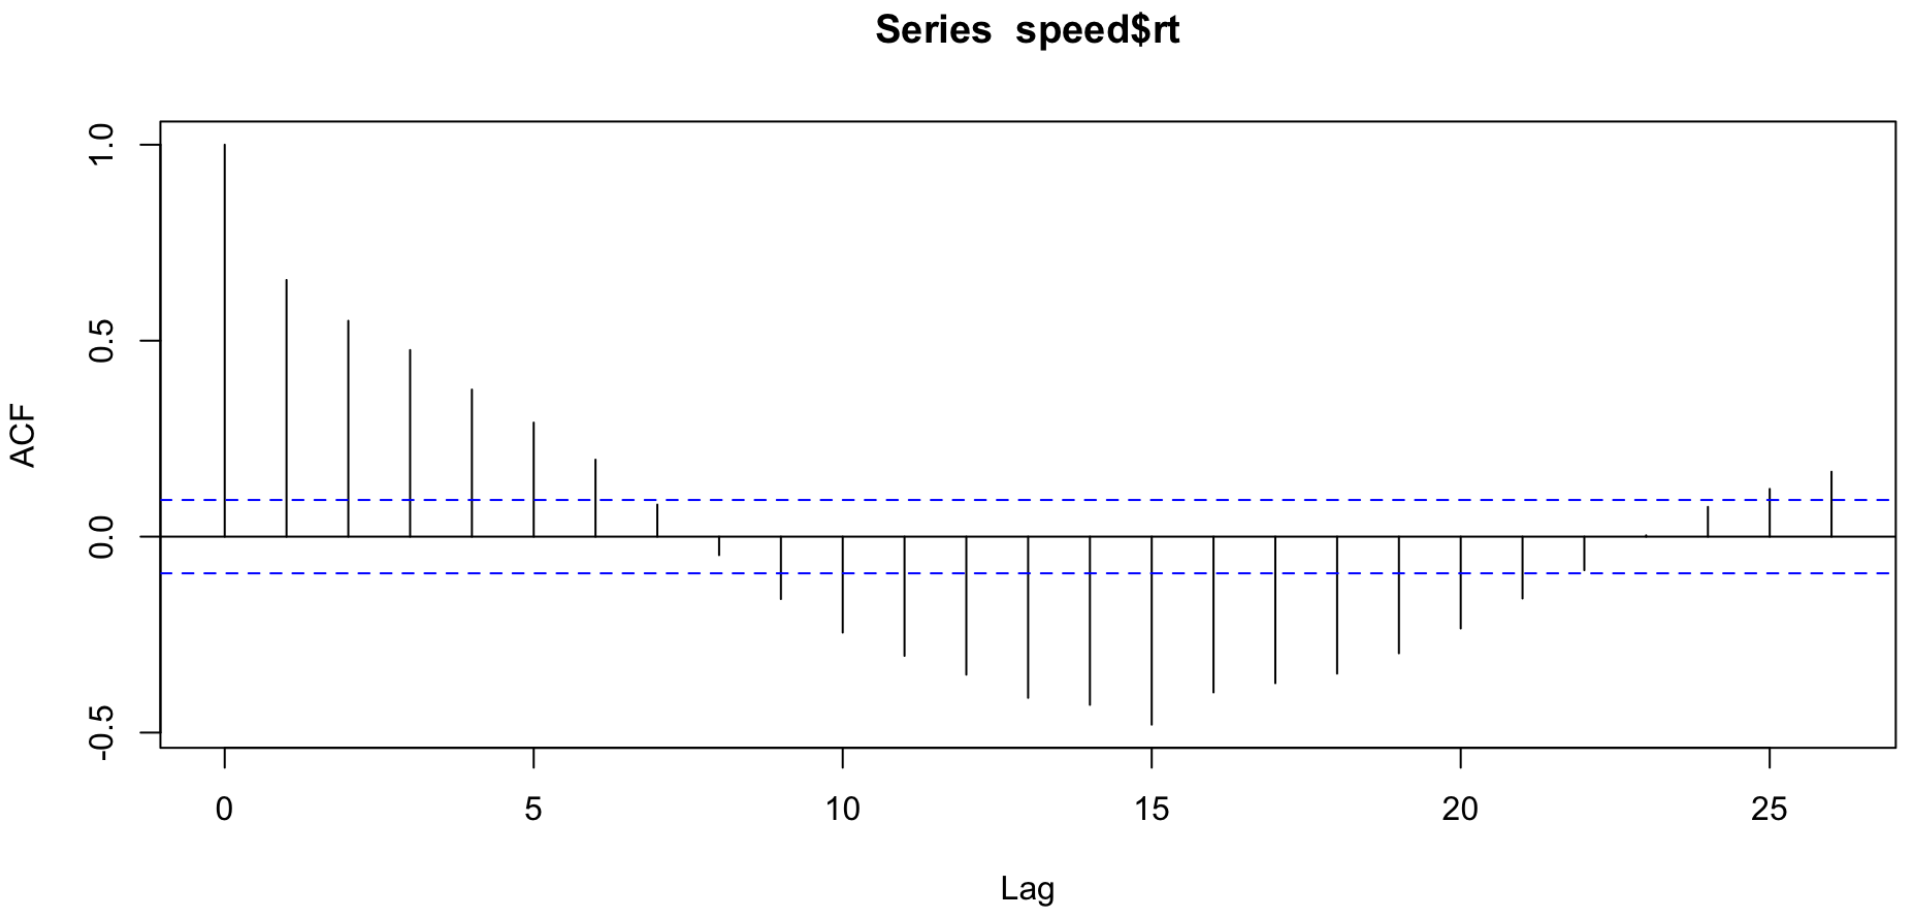

A simple hidden Markov model

```
data(speed)
mod1 <- depmix(list(rt~1,corr~1),
               data=speed,
               nstates=2,
               family=list(gaussian(),multinomial("identity")),
               ntimes=c(168,134,137))
# fit the model
set.seed(1234)
fmod1 <- fit(mod1, verbose=FALSE)
```

```
converged at iteration 24 with logLik: -296.1
```

Parameter estimates

```
Initial state probabilities model
pr1 pr2
  0   1

Transition matrix
        toS1  toS2
fromS1 0.899 0.101
fromS2 0.084 0.916

Response parameters
Resp 1 : gaussian
Resp 2 : multinomial
    Re1.(Intercept) Re1.sd Re2.inc Re2.cor
St1           5.521  0.202   0.472   0.528
St2           6.392  0.240   0.099   0.901
```

Accounting for the effect of Pacc

Using Pacc as a covariate on the transition probabilities

```
data(speed)
mod2 <- depmix(list(rt~1,corr~1),
               data=speed,
               transition=~Pacc,
               nstates=2,
               family=list(gaussian(),multinomial("identity")),
               ntimes=c(168,134,137))
# fit the model
set.seed(1234)
fmod2 <- fit(mod2, verbose=FALSE)
```

```
converged at iteration 24 with logLik: -249
```

Covariate on transition probabilities

Transition probabilities are modeled using multinomial logistic regression

$$\text{logit}(1 - a_{11}(t)) = \eta_0^1 + \eta_1^1 \cdot \texttt{Pacc}_t.S$$

$$\text{logit}(a_{22}(t)) = \eta_0^2 + \eta_1^2 \cdot \texttt{Pacc}_t.$$

Model parameters

```
Response parameters
Resp 1 : gaussian
Resp 2 : multinomial
    Re1.(Intercept) Re1.sd Re2.inc Re2.cor
St1           5.522  0.203   0.474   0.526
St2           6.394  0.237   0.096   0.904
```

Model parameters

```
Transition model for state (component) 1
Model of type multinomial (mlogit), formula: ~Pacc
Coefficients:
            St1    St2
(Intercept)   0 -4.223
Pacc          0  9.134
Probalities at zero values of the covariates.
0.9856 0.01445

Transition model for state (component) 2
Model of type multinomial (mlogit), formula: ~Pacc
Coefficients:
            St1    St2
(Intercept)   0 -3.373
Pacc          0 15.804
Probalities at zero values of the covariates.
0.9669 0.03314
```

Transition probability function

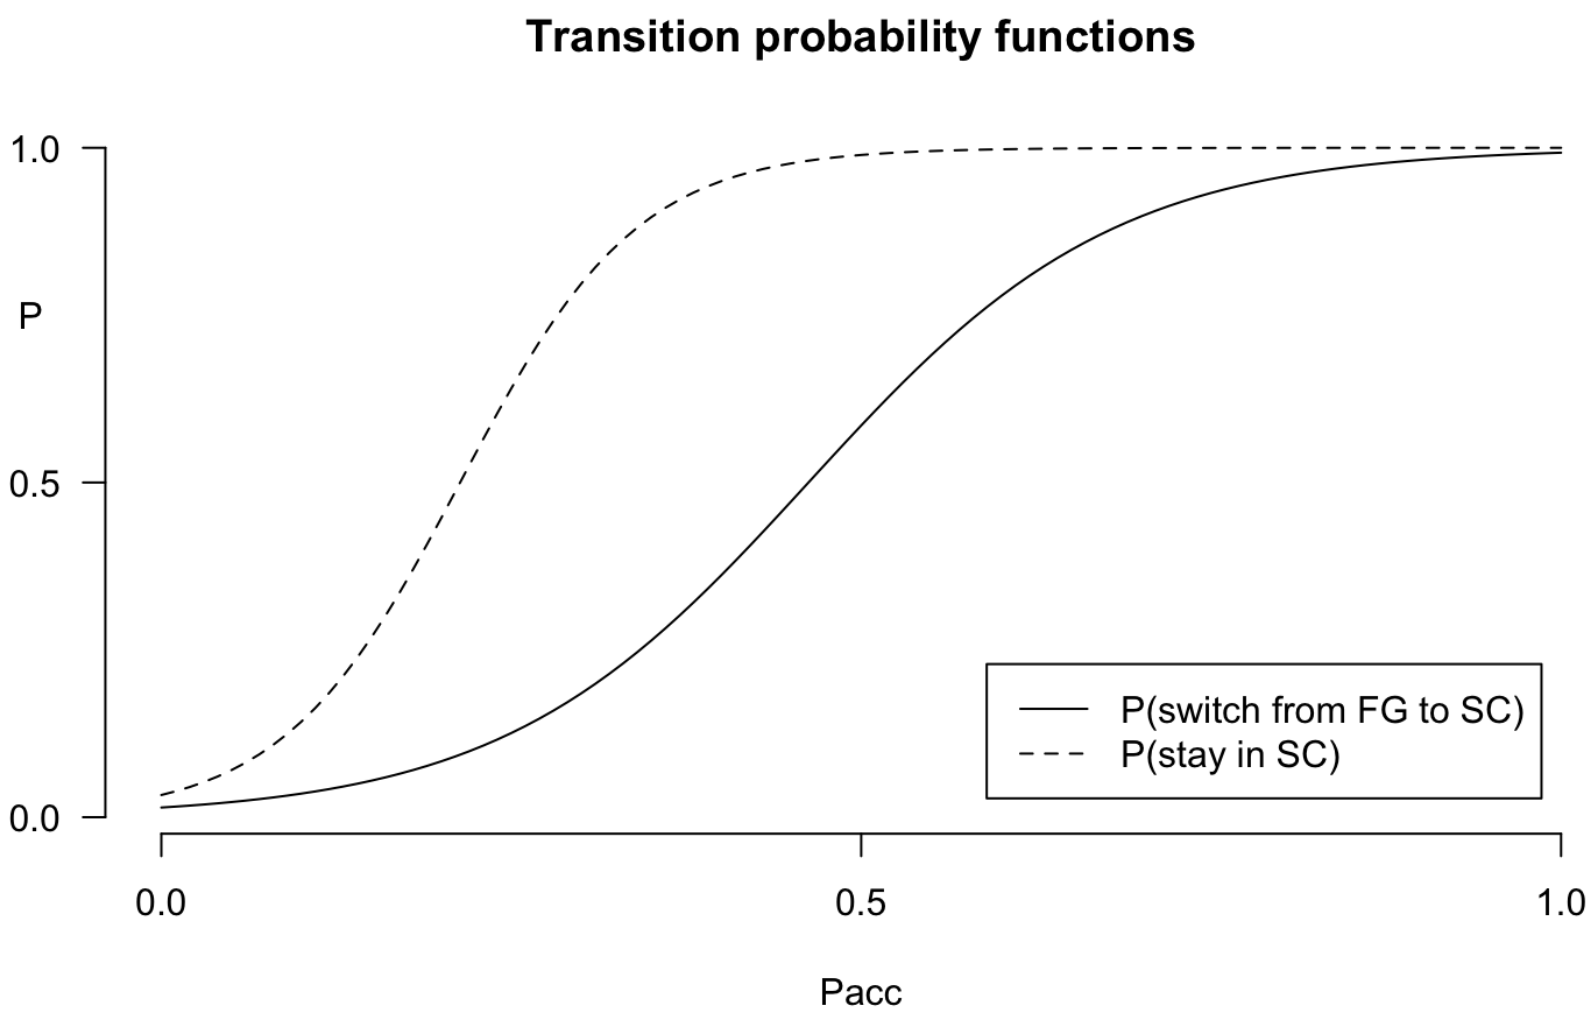

Model checking

Autocorrelation function of the residuals of the RTs

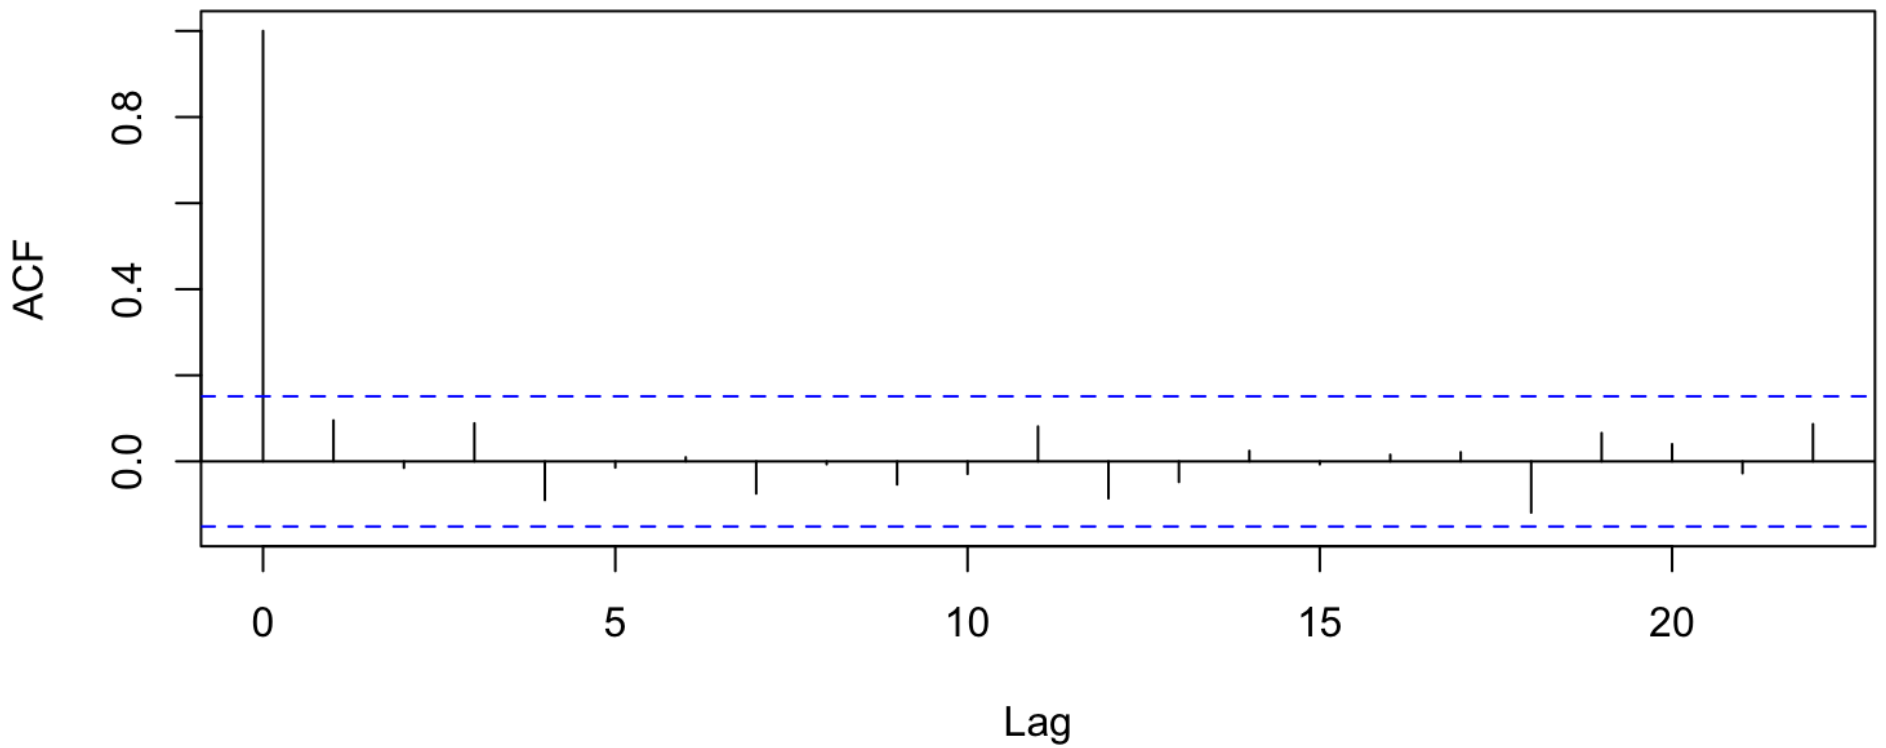

Parameter constraints and inference

```
pars <- c(unlist(getpars(fmod2)))

# constrain the initial state probs to be 0 and 1
# also constrain the guessing probs to be 0.5 and 0.5
# (ie the probabilities of corr in state 1)
pars[1]=0
pars[2]=1 # this means the process will always start in state 2
pars[13]=0.5
pars[14]=0.5
mod3 <- setpars(mod2,pars)
```

Parameter constraints and inference

```
1 summary(mod3,which="prior")
```

```
Initial state probabilities model
pr1 pr2
  0   1
```

```
1 summary(mod3,which="response")
```

```
Response parameters
Resp 1 : gaussian
Resp 2 : multinomial
    Re1.(Intercept) Re1.sd Re2.inc Re2.cor
St1           5.522  0.203   0.500   0.500
St2           6.394  0.237   0.096   0.904
```

Parameter constraints and inference

```r
# fix the parameters by setting:
free <- c(0,0,rep(c(0,1),4),1,1,0,0,1,1,1,1)
# fit the model
fmod3 <- fit(mod3,fixed=!free)
```

```
Iter: 1 fn: 249.2129     Pars:  -4.22283  9.13384 -3.37335 15.80431  5.52172
0.20292  6.39370  0.23736  0.09572  0.90428
solnp--> Completed in 1 iterations
```

Parameter constraints and inference

Likelihood ratio test on the fitted models

```
1 fmod2
```

```
Convergence info: Log likelihood converged to within tol. (relative change)
'log Lik.' -249 (df=11)
AIC:  519.9
BIC:  564.9
```

```
1 fmod3
```

```
Convergence info: 0  (0 is good in Rsolnp, check manual for other values)
'log Lik.' -249.2 (df=7)
AIC:  512.4
BIC:  541
```

```
1 # likelihood ratio insignificant, hence fmod3 better than fmod2
2 llratio(fmod2,fmod3)
```

```
log Likelihood ratio (chi^2): 0.481 (df=4), p=0.975.
```

Discussion

- Studying direct vs indirect effects
-