



Prediction model for travel time variability

Fosgerau, Mogens; Hjorth, Katrine; Jensen, Thomas Christian; Prameswari, Niken

Publication date:
2015

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Fosgerau, M., Hjorth, K., Jensen, T. C., & Prameswari, N. (2015). *Prediction model for travel time variability*. DTU Transport.

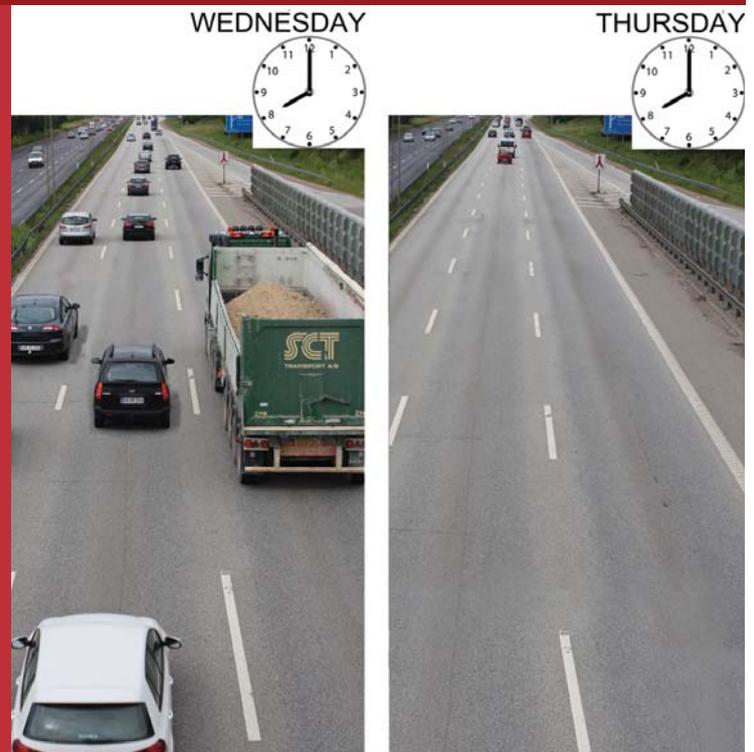
General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Prediction model for travel time variability



Mogens Fosgerau
Katrine Hjorth
Thomas C. Jensen
Niken Prameswari

June 2015

Prediction model for travel time variability

Mogens Fosgerau

Katrine Hjorth

Thomas C. Jensen

Niken Prameswari

2015

Prediction model for travel time variability

Report 10
2015

By
Mogens Fosgerau
Katrine Hjorth
Thomas C. Jensen
Niken Prameswari

Copyright: Reproduction of this publication in whole or in part must include the customary bibliographic citation, including author attribution, report title, etc.

Cover photo: DTU Transport

Published by: Department of Transport, Bygningstorvet 116B

Request report

from: www.dtu.dk

ISSN: 1601-9458 (electronic version)

ISBN: 87-7327-281-7 (electronic version)

Content

Summary in Danish (dansk sammenfatning)	i
Projektets formål	i
Datagrundlag	ii
Metoden	iii
Eksempler på anvendelser	iv
Referencer	x
1. Introduction	1
2. Data description	3
2.1 Choice of main data source	3
2.2 Additional data sources used in analysis	4
2.3 Definition of analysis variables	5
2.4 Sample selection	8
2.5 Descriptive statistics	9
3. Model	19
3.1 Central concepts and assumptions	20
3.2 Breakdown Model	20
3.3 Recovery Model	21
3.4 Travel Time Model	22
4. Empirical analysis	24
4.1 Definition of link types	24
4.2 Definition of breakdown and recovery times	24
4.3 Estimation of breakdown model	26
4.4 Estimation of recovery model	28
4.5 Model validation	32
4.6 Estimated state-dependent means and variances of travel time	34
5. Application	35
5.1 Outline of implementation	35
5.2 Application example	36
5.3 Applying the model for experiments	39
6. Conclusions and recommendations for future model development, implementation and use in traffic model	46
6.1 Handling spillback effects	46
6.2 Generalising to other motorways and remaining road network	46
References	48
Appendix A Tables	49
Acknowledgements	52

Summary in Danish (dansk sammenfatning)

Projektets formål

Den stigende mængde trafik på vejene giver mere udbredt trængsel, som medfører dels en stigning i de gennemsnitlige rejsetider, dels at de enkelte rejsetider i stigende grad bliver variable og uforudsigelige. Denne uforudsigelighed kaldes rejsetidsvariabilitet (TTV, for *travel time variability*) og har potentielt store samfundsmæssige omkostninger. Der findes dog ikke på nuværende tidspunkt en veletableret praksis for, hvordan ændringer i TTV skal opgøres i samfundsøkonomiske projektvurderinger. Foreløbige danske regneeksempler anslår, at man i den samfundsøkonomiske analyse undervurderer rejsetidsomkostningerne med 10-20% på steder med meget rejsetidsvariabilitet, hvis man ikke inkluderer TTV. Det kan have stor betydning for det samlede resultat af analysen, idet sparede rejsetidsomkostninger som regel udgør 60-80% af gevinsten ved infrastrukturprojekter (DTU Transport, 2008).

Den nuværende praksis tager til en vis grad højde for usikkerhed i rejsetider, ider der skelnes mellem fri køretid og gennemsnitlige forsinkelser i forhold til denne (dette gælder for privat transport, mens der for kollektiv transport skelnes mellem køreplanstid og gennemsnitlige forsinkelser). Gennemsnitlige forsinkelser vægtes med hhv. 1,5 for passagerbiler og 1,4 for last-og varebiler (og 2,0 for kollektiv transport) i beregningen af rejsetidsomkostninger i samfundsøkonomiske analyser. Denne metode tager dog ikke højde for omfanget af uforudseede forsinkelser. Af denne grund er der en risiko for, at den nuværende praksis undervurderer de egentlige omkostninger af trængsel. Det anbefales derfor (DTU Transport, 2008) at skifte til en opgørelsesmetode, hvor man – i stedet for at skelne mellem fri køretid/køreplanstid og gennemsnitlige forsinkelser - skelner mellem gennemsnitlig rejsetid og TTV, da der er fagligt grundlag for at tillægge disse størrelser en samfundsøkonomisk værdi. Dette skifte kræver dog to ting: For det første skal man kende den samfundsøkonomiske værdi af TTV, og for det andet skal man kunne opgøre niveauet af TTV i trafikscenarier.

Transportministeriet og Vejdirektoratet har derfor bedt DTU Transport udvikle en praktisk anvendelig metode, der kan bruges til at forudsige niveauet af TTV i trafikprognoser for vejnettet. Denne rapport beskriver metoden og giver eksempler på dens anvendelse. Den udviklede model er tænkt som et efterberegningsmodul til Landstrafikmodellen (LTM) eller en alternativ trafikmodel, der forudsiger trafikmængder fordelt over døgnets forskellige tidsintervaller. Det er afgørende, at tid på døgnet indgår i trafikmodellen, da trafikfordelingen over døgnet selvsagt har afgørende betydning for niveauet af trængsel. Anvendelse af metoden kræver derfor LTM's version 2.0.

Med den nye metode er vi i stand til at måle, hvordan omfanget af TTV på motorvejen påvirkes af ændrede trafikmønstre eller ændringer i antallet af spor. I dette sammendrag viser vi fire simple eksempler, der illustrerer metodens anvendelsesmuligheder:

- En ændring i trafikmønstret, således at en del af bilisterne flyttes fra myldretiden til perioderne før og efter (f.eks. som følge af road pricing).
- Indførsel af et intelligent rampedoseringsystem, der løbende justerer trafiktilførslen til motorvejen via sluser på ramperne alt efter trafiksituationen. I illustrationen er dette modelleret ved at sandsynligheden for at trængsel opstår nedbringes med 20%.
- En udvidelse af en strækning fra tre til fire spor (her forudsættes dog, at trafikefterspørgslen ikke ændres, da modellering af dette kræver anvendelse af LTM)
- En udvidelse af en strækning fra to til tre spor (her forudsættes dog, at trafikefterspørgslen ikke ændres, da modellering af dette kræver anvendelse af LTM)

Som det diskuteres nedenfor er modellen en prototype, der er estimeret på data fra Køge Bugt Motorvejen, og ideelt set bør udvides, så den baseres på flere motorvejsstrækninger samt andre vejtyper. Dette er kun muligt i det omfang de nødvendige data er tilgængelige, hvilket i øjeblikket kun gælder for visse dele af motorvejsnettet. Modellen bør tillige forsøges udvidet, så den tager højde for spillback-effekter (også kaldet tilbagestuvning), hvilket ikke er opnået i løbet af projektet. Det skønnes muligt i det mindste at kunne approksimere effekten af spillbacks med de anvendte data, mens en detaljeret analyse af flaskehalse og spillbacks kræver langt mere detaljerede data.

Når TTV (målt som rejsetidens standardafvigelse) skal omregnes til generaliserede rejsetidsomkostninger i samfundsøkonomiske analyser, skal der anvendes en samfundsøkonomisk værdi for TTV. Den anbefalede værdi for et minuts standardafvigelse er pt. lig tidsværdien for et minuts rejsetid (jf. DTU Transport, 2008). Denne værdi er baseret på et review af internationale studier og forventes revideret i løbet af den nærmeste fremtid i forbindelse med nye internationale erfaringer og et nyt dansk forskningsstudie på DTU Transport.

Datagrundlag

Den udviklede metode er en model, der forudsiger niveauet af TTV for en given trafikprofil (trafikmængder over døgnets forskellige tidsintervaller) for danske motorveje. Som nævnt er modellen kalibreret på data fra Køge Bugt Motorvejen, og bør udvides til også at dække flere vejtyper, når de nødvendige data er tilgængelige. Da datakravene er betydelige (observationer af rejsetid og trafikmængder målt over sammenhængende tidsintervaller over mange dage) er det imidlertid ikke realistisk på kort sigt at udvikle separate modeller for alle vejtyper. Vi anbefaler derfor, at man prioriterer at kalibrere separate modeller for forskellige typer af motorvejsstrækninger (flere end i denne rapport) samt andre større veje, hvor datagrundlaget allerede er til stede. For visse motorvejsstrækninger, f.eks. Helsingør Motorvejen, er de relevante data tilgængelige. Det samme er muligvis gældende for enkelte af de øvrige større veje, mens datagrundlaget for eksempelvis de kommunale veje skønnes utilstrækkeligt.¹

Det er vigtigt at være opmærksom på, hvorvidt givne datakilder indeholder tilstrækkeligt detaljeret information til at kunne bruges til at kalibrere modellen. Dette gælder særligt i forbindelse med fremtidige dataindsamlinger, der sættes i gang. Først og fremmest er det nødvendigt at have målinger af rejsetid og trafikflow både før, under og efter myldretiden for at kunne identificere, hvornår der opstår trængsel, og analysere den dynamiske proces der foregår under afviklingen af trængsel. Rejsetider på strækningsniveau er at foretrække (frem for på målepunktniveau, som er anvendt i nærværende analyse i mangel af bedre). For at kunne beregne niveauet af TTV på et tidspunkt med en vis sikkerhed, er det desuden nødvendigt med gentagne målinger af rejsetiden på dette tidspunkt, både på en given dag og over mange dage.

Det skal understreges, at man kun kan kalibrere modellen, hvis der er trængsel på den pågældende vejstrækning. Hvis der ikke er trængsel med den nuværende trafikmængde, kan man ikke identificere, hvilke trafikmængder, der skal til, før der opstår trængsel. Har man brug for at forecaste niveauet af TTV for en sådan vejstrækning (i et scenarie hvor der forventes at opstå trængsel), skal der anvendes modeller kalibreret på tilsvarende vejstrækninger.

I forbindelse med en mere detaljeret modellering af spillbacks er det desuden af stor betydning, at man observerer alle relevante trafikflows hen mod og væk fra en flaskehals, dvs. det er gavnligt så vidt muligt at måle trafikken på til- og frakørselsramper. Sådanne data var ikke tilgængelige da modellen blev udviklet.

¹ Vi henviser til, at Vejdirektoratet, i forbindelse med et andet projekt om udvikling af indikatorer for trængsel, forventes at udarbejde et notat med overblik over eksisterende datakilder og deres omfang i foråret 2015.

Vi henviser til rapportens sektion 2 for yderligere information om de anvendte data og udvælgelsesprocessen.

Metoden

Metoden består af en statistisk model, kalibreret på data fra Køge Bugt Motorvejen, og en simulationsmodel. Den statistiske model beskriver, hvordan sandsynligheden for at trængsel opstår i et givet tidsinterval og sandsynligheden for at trængslen afvikles igen i et givet tidsinterval afhænger af trafikflowet (antal biler pr. spor pr. minut) i døgnet's tidsintervaller. Simulationsmodellen anvender den statistiske model til at simulere rejsetidens middelværdi og standardafvigelse for hvert enkelt tidsinterval, baseret på trafikflows estimeret i LTM.

Den statistiske model

De vigtigste principper i vores arbejde med at udvikle den statistiske model er:

- Modellen skal kunne forudsige TTV med rimelig sikkerhed på et **aggregeret niveau**. Dvs. vi søger at modellere et overordnet forhold mellem TTV og trafikmængden, som kan siges at være generelt gældende for alle motorvejsstrækninger. Det viser sig (ikke overraskende), at forholdet mellem rejsetid og trafikmængde varierer meget indenfor de forskellige vejstrækninger i analysen, fordi rejsetid og trængsel afhænger af den enkelte vejstrækning's udformning, som afgør hvor flaskehalsene opstår og hvor store deres konsekvenser er. En model som er i stand til detaljeret at forudsige rejsetid og TTV for hvert enkelt vejstrækning kræver derfor detaljeret information om strækningens udformning, samt (potentielt) separat statistisk modellering af hver enkelt strækning for sig. Det er i teorien muligt at opstille en sådan model, men det skønnes at være et ganske omfattende arbejde, som i princippet hører under den fremtidige udvikling af Landstrafikmodellen. Det teoretiske grundlag for modellering af flaskehalse er desuden et område, der er under stadig udvikling i forskningsstudier, bl.a. Ph.D.-projekter på DTU. For at nå frem til en metode, som kan anvendes allerede sammen med LTM's version 2.0, er det derfor nødvendigt at fokusere på at modellere et mere aggregeret niveau.
- Sammenhængen mellem trafikflow og rejsetid beskrives traditionelt vha. speed-flow kurver (f.eks. i LTM). **Et væsentligt princip i den nye metode er, at den tager højde for to vigtige teoretiske problemstillinger, som traditionelle speed-flow kurver ikke tager højde for.** Det drejer sig om:
 - o Når der er trængsel, afhænger det målte trafikflow af hastigheden, fordi mængden af biler, der passerer et givent målepunkt pr. tidsenhed afhænger af hastigheden. Det betyder, at trafikflowet er en endogen variabel i forhold til rejsetiden (hvilket vil sige, at variabelen ikke kan antages at være ukorreleret med fejleddet i en statistisk model). De eksisterende speed-flow kurver tager ikke højde for dette, og der er derfor stor risiko for systematiske fejl i de estimerede parametre (dvs. at kurverne er misvisende).
 - o Trængsel er et dynamisk fænomen, der afvikles over tid. I trængselsperioderne vil rejsetiderne i de enkelte tidsperioder typisk være indbyrdes afhængige, fordi det tager tid at afvikle trængsel igen, når den først er opstået. De eksisterende speed-flow kurver tager ikke højde for dette, hvilket også kan medføre systematiske fejl.
- En anden vigtig problemstilling er spillback-effekter, altså hvordan rejsetiden på en strækning påvirkes af trafikforholdene længere fremme. Dette er også noget, som traditionelle speed-flow kurver ikke tager højde for. I projektoplægget lægges der op til, at den nye metode skal tage højde

for spillback-effekter, men dette har vi *ikke* opnået med den nuværende modelformulering. Vi anbefaler dog, at dette undersøges nærmere i forbindelse med videre udvikling og re-estimation af modellen.

Den statistiske model er en simplificeret beskrivelse af virkeligheden. Vi antager, at trafikken har to tilstande: Ikke-trængsel og trængsel. Vi betragter en periode (f.eks. morgenperioden 5:00-12:00 eller eftermiddagsperioden 12:00-19:00), hvori vi antager trafikken kan skifte tilstand fra ikke-trængsel til trængsel og tilbage igen højst én gang. Modellen består af to delmodeller:

- Breakdownmodellen: Bestemmer sandsynligheden for, at trængsel opstår, for hvert 15-minutters tidsinterval, givet at trafikken stadig er i ikke-trængsels-tilstand. Sandsynligheden er modelleret som en simpel logistisk funktion af trafikflowet i tidsintervallet. Idet der endnu ikke er opstået trængsel, kan vi antage, at trafikflowet ikke er endogent i modellen.
- Recoverymodellen: Bestemmer sandsynligheden for, at trængslen slutter, for hvert 15-minutters tidsinterval, givet at trafikken stadig er i trængsels-tilstand. Sandsynligheden er modelleret som en simpel logistisk funktion af det gennemsnitlige flow siden starten på trængselsperioden, en formulering med inspiration i flaskehalsmodellen i den teoretiske trafiklitteratur (de Palma and Fosgerau, 2011).

Vi henviser til rapportens sektion 3 og 4 for yderligere information.

Simulationsmodellen

Inputtet i simulationsmodellen er en trafikefterspørgselsprofil (trafikefterspørgsel fordelt på et antal tidsbånd) for en typisk hverdag. Simulationsmodellen er en algoritme, der først genererer et stort antal dage med variable efterspørgselsprofiler, der gennemsnitligt svarer til inputprofilen. For hver dag simuleres, hvorvidt og hvornår der opstår trængsel, og hvornår den i så fald slutter igen – her bruges den statistiske model. Når man betragter alle de simulerede dage under et, kan man for hvert 15-minutters tidsinterval beregne sandsynligheden for, at trafikken er i hhv. Ikke-trængsel og trængsel, samt den gennemsnitlige rejsetid og TTV.

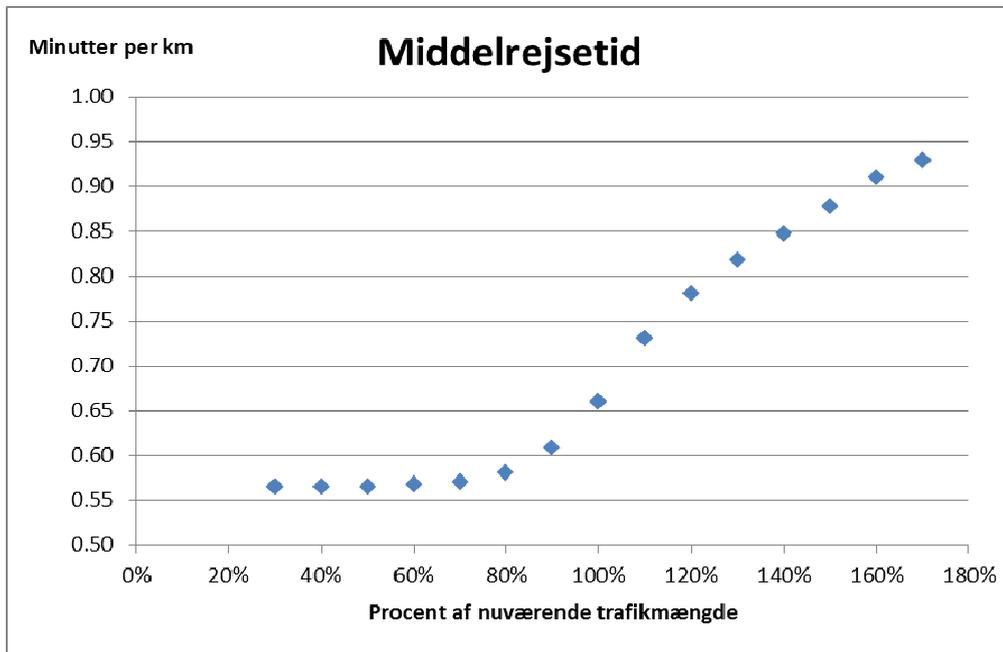
Vi henviser til rapportens sektion 5 for yderligere information.

Eksempler på anvendelser

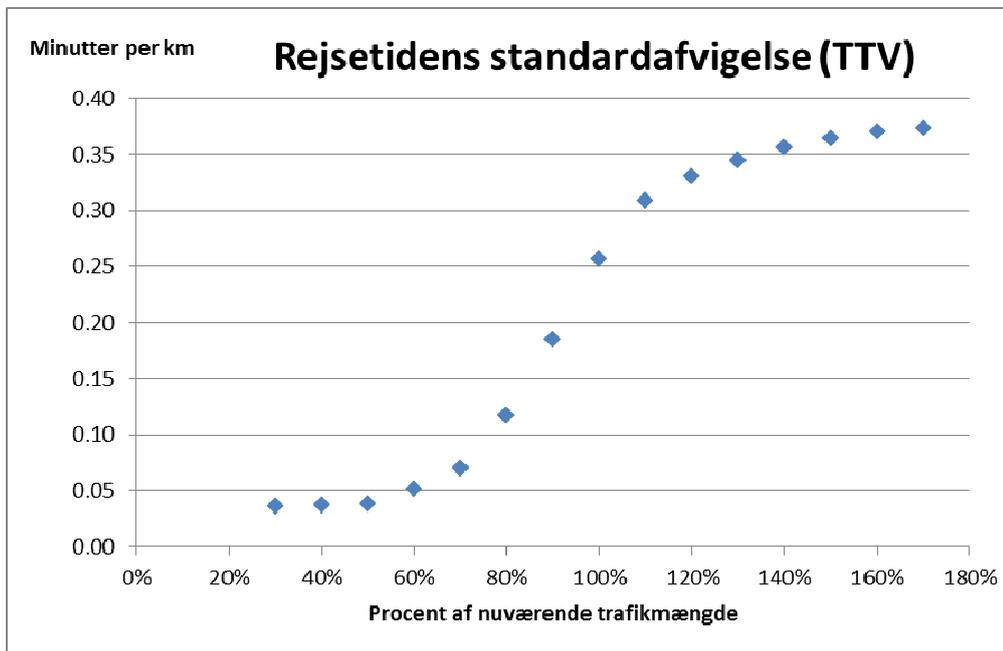
Vores første eksempel på en anvendelse af modellen er simulation af såkaldte volume-delay sammenhænge, altså sammenhænge mellem gennemsnitlig rejsetid og trafikvolumen hhv. TTV og trafikvolumen. Vi beregner disse sammenhænge ved at skalere den nuværende trafikefterspørgsel med 30%, 40%, 50% 170% og simulere gennemsnitsrejsetiden og TTV for hvert scenarie. *Figur 1* og *Figur 2* viser de resulterende gennemsnitsrejsetider og TTV (vist som vægtede gennemsnit over hele morgenperioden 4:30-12:00, vægtet med trafikmængden i hvert tidsinterval). I figurerne er brugt 1000 gentagelser af hvert scenarie.

Som forventet stiger både gennemsnitsrejsetiden og TTV med trafikvolumen. Ved meget lave trafikmængder bliver de konstante. Det skyldes, at sandsynligheden for trængsel bliver så lille, at det (mere eller mindre) aldrig indtræffer. Ved meget høje trafikmængder sker det modsatte: I de ekstreme tilfælde bliver trafikmængden så stor, at der er trængsel i alle tidsintervaller, og modellen antager så, at den høje (men konstante) gennemsnitsrejsetid og TTV, vi observerer i trængselstilstand, gælder for alle tidsperioder. Modellen er derfor ikke helt realistisk ved meget høje trafikniveauer: Den opererer kun med to

tilstande (trængsel og ikke-trængsel), der hver antages at have en konstant fordeling af rejsetider uafhængigt af trafikmængden. Vi forventer, at modellen undervurderer både gennemsnitsrejsetid og TTV ved meget høje trafikmængder.



Figur 1: Simuleret gns. rejsetid (gennemsnit over perioden 4:30-12:00) som funktion af trafikvolumen.



Figur 2: Simuleret standardafvigelse af rejsetiden (gennemsnit over perioden 4:30-12:00) som funktion af trafikvolumen.

Vi har også anvendt modellen til at simulere effekten af en række trængselsreducerende tiltag. Nedenstående eksempler er tænkt som en illustration af metodens anvendelsesmuligheder. Det er højst simplificerede case-scenarier, som ikke nødvendigvis er 100% realistiske. F.eks. antages i alle tilfælde, at den samlede trafikefterspørgsel er uændret. Mere realistiske beregninger kræver, at trafikefterspørgslen først beregnes i LTM.

Cases I-III er beregnet for en motorvejsstrækning med tre spor, mens case IV er beregnet for en motorvejsstrækning med to spor.

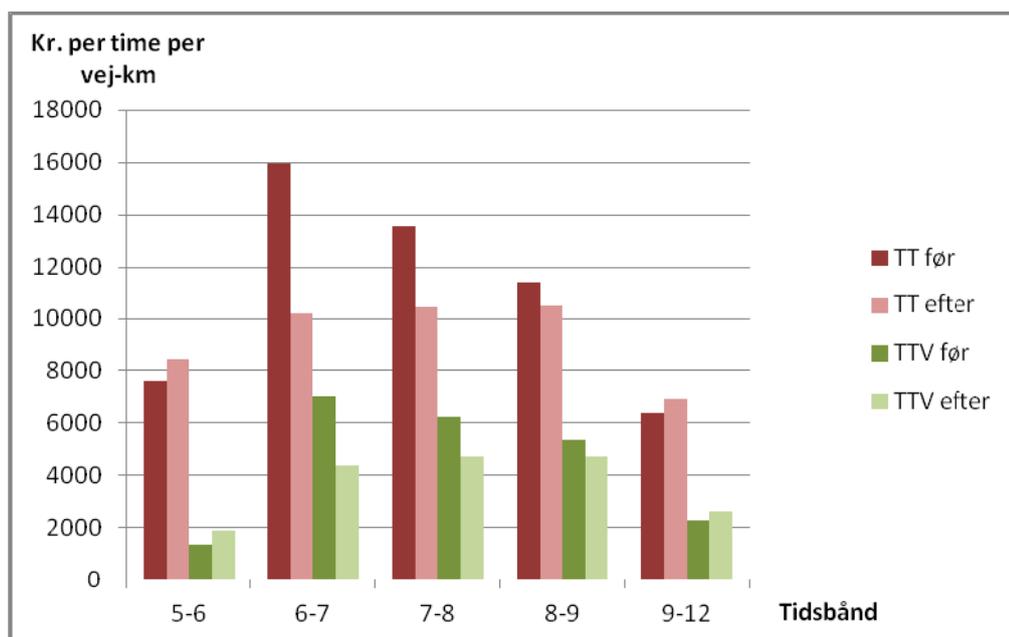
Den anvendte tidsværdi for et minuts gennemsnitlig rejsetid er vægtede gennemsnit af de gældende tidsværdier for personbiler, varebiler og lastbiler, hvor vægtningen varierer over dagen og antages at være den samme som i basisscenariet.

Vi henviser til rapportens sektion 5 for yderligere information om case beregningerne.

Case I: Peak spreading med road pricing

Denne case er modelleret ved at trafikefterspørgslen jævnes ud, så det aldrig overstiger 25 personbilsækvivalenter pr. spor pr. minut. "Overskydende" efterspørgsel flyttes så lidt som muligt til tidsperioder før og efter: Halvdelen flyttes til før den oprindelige myldretid, halvdelen til efter. Mere specifikt betyder det, at den gennemsnitlige efterspørgsel nedjusteres i perioden kl. 5:30-8:30, og opjusteres i perioderne kl. 5:00-5:30 samt kl. 8:30-9:45. Dette er en simpel imitation af konsekvenserne af en tidsafhængige kørselsafgifter, der er højere i den oprindelige myldretidsperiode end udenfor.

Figur 3 sammenligner rejsetidsomkostningerne i case I med basisscenariet. Omkostningerne er fordelt på værdien af den gennemsnitlige rejsetid (røde søjler) og værdien af TTV (grønne søjler). I perioderne uden for den oprindelige myldretid (kl. 5-6 og kl. 9-12) sker der en stigning i omkostningerne, fordi både gennemsnitlig rejsetid og TTV stiger som følge af den ekstra trafik. Stigningen mere end opvejes dog af gevinsten i perioden kl. 6-9, hvor der sker betydelige fald i omkostningerne. Dette er fuldt ud forventeligt i dette scenarie. Det væsentlige er her, at omkostninger fra rejsetidsvariabilitet udgør en signifikant del af de samlede omkostninger, jf. Tabel 1. Medregnes disse omkostninger i den samfundsøkonomiske analyse, skønnes de derfor at have stor betydning for resultatet. Samlet set falder rejsetidsomkostningerne med ca. 11%, jf. Tabel 1.

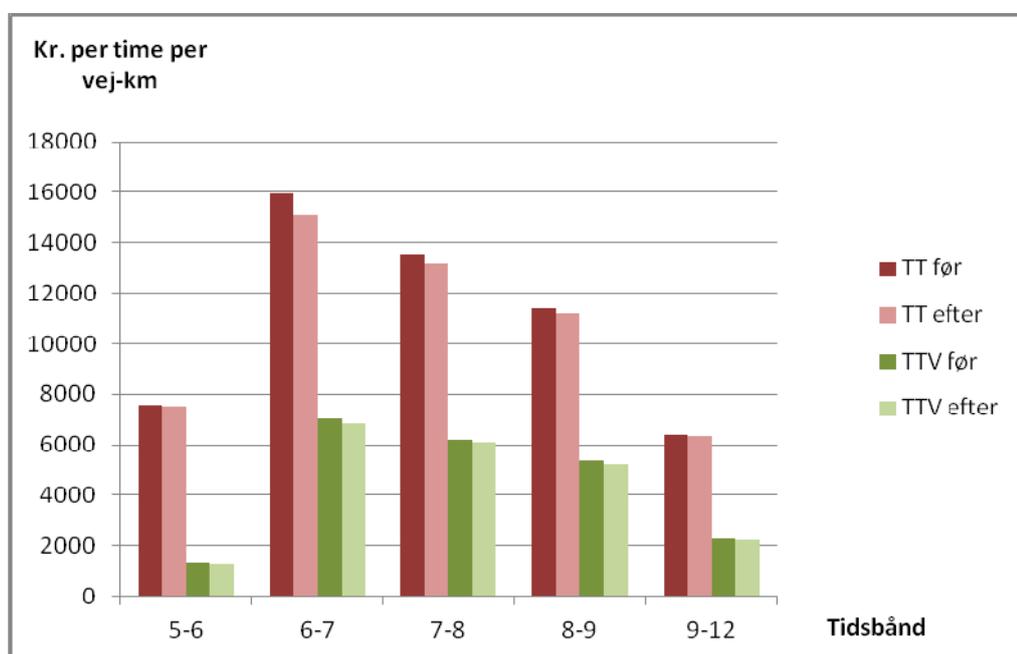


Figur 3: Værdien af gns. rejsetid (TT) og rejsetidsvariabilitet (TTV) i basisscenariet og case I.

Case II: Rampedoseringssystem

Denne case er modelleret ved, at sandsynligheden for at trængsel opstår, nedbringes med 20% på ethvert tidspunkt på dagen. Dette er en simpel imitation af et rampedoseringssystem der kontrollerer trafikmængden på motorvejen ved at tilbageholde biler på ramperne i kortere perioder. De negative effekter fra tilbageholdte biler er ikke medregnet. Eventuelle ændringer i trafikefterspørgslen er ligeledes ikke medregnet.

Figur 4 viser rejsetidsomkostningerne for case II og sammenligner med basisscenariet. Case II giver en begrænset besparelse som følge af en reduktion af gennemsnitlige rejsetid, mens ændringerne i niveauet af TTV er meget små. De meget små ændringer skyldes formentlig, at en ændring i sandsynligheden på 20% er så lille, at den generelt blot udskyder myldretiden en smule og overordnet set kun sænker andelen af dage med trængsel fra ca. 89% til 83%.

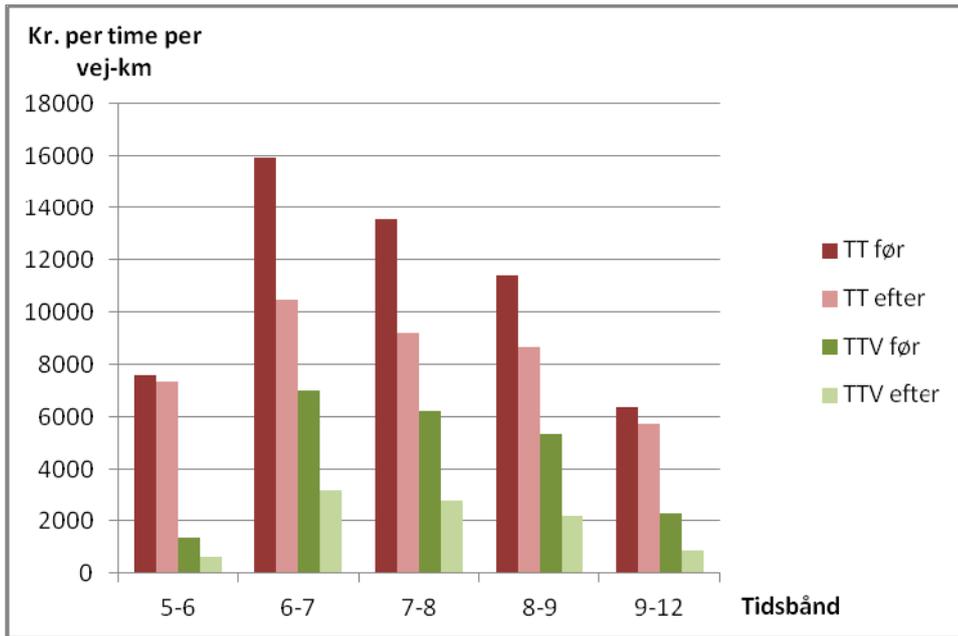


Figur 4: Værdien af gns. rejsetid (TT) og rejsetidsvariabilitet (TTV) i basisscenariet og case II.

Case III: Udvidelse fra tre til fire spor

Udvidelsen fra tre til fire spor er modelleret ved, at den samlede nuværende trafikmængde på en tresporet motorvejsstrækning simpelthen fordeles på fire spor i stedet for tre. Vi understreger, at dette er en højst forsimplet antagelse: En mere realistisk modellering kræver, at ændringer i trafikefterspørgslen simuleres i LTM.

Figur 5 viser rejsetidsomkostningerne for case III og sammenligner med basisscenariet. Case III giver markante forbedringer i gennemsnitlig rejsetid inden for myldretiden, og markante forbedringer i TTV i alle perioder. Igen er omkostningerne forbundet med den sparede rejsetidsvariabilitet tilstrækkeligt store til at påvirke resultatet i en samfundsøkonomisk analyse. De samlede rejsetidsomkostninger falder med ca. 32%, jf. Tabel 1.

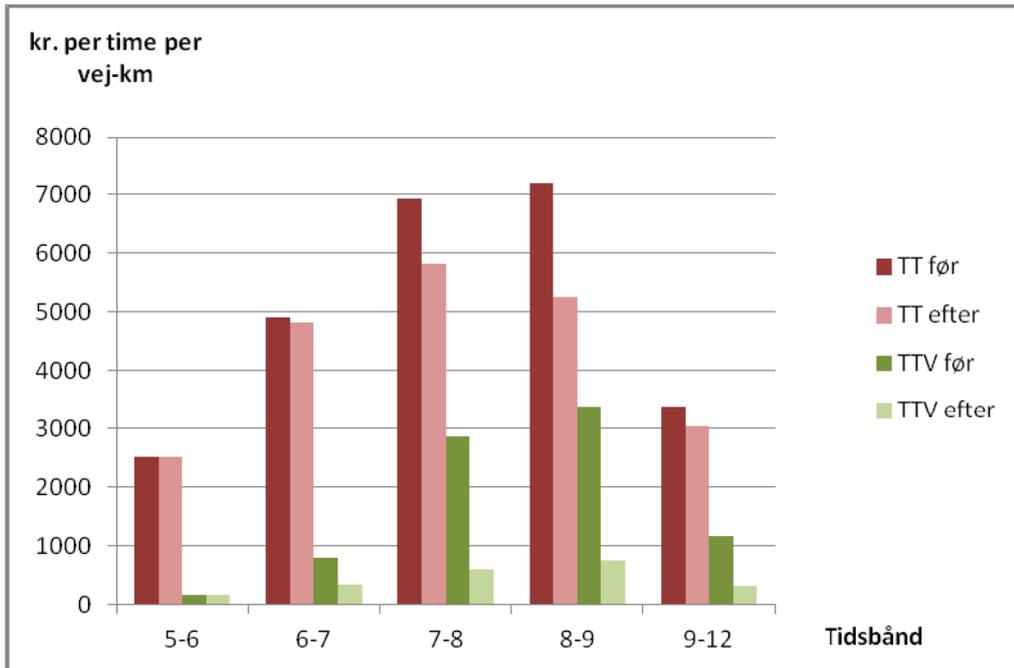


Figur 5: Værdien af gns. rejsetid (TT) og rejsetidsvariabilitet (TTV) i basisscenariet og case III.

Case IV: Udvidelse fra to til tre spor

Udvidelsen fra to til tre spor er modelleret ved, at den samlede nuværende trafikmængde på en tosporet motorvejsstrækning fordeles på tre spor i stedet for to. Der gælder derfor samme bemærkning som ovenfor: En mere realistisk modellering kræver, at ændringer i trafikefterspørgslen først simuleres i LTM. Beregningen er yderligere forenklet ved at antage, at hele strækningen i basisscenariet kun har to spor, mens den i virkeligheden har tre spor på en del af strækningen.

Figur 6 viser rejsetidsomkostningerne for case IV og sammenligner med basisscenariet. Case IV giver markante forbedringer i både gennemsnitlig rejsetid og TTV inden for myldretiden, som for denne strækning er kl. 7-9, samt en relativt stor forbedring TTV efter myldretiden (kl. 9-12). Igen er omkostningerne forbundet med den sparede TTV tilstrækkeligt store til at påvirke resultatet i en samfundsøkonomisk analyse. De samlede rejsetidsomkostninger falder med ca. 28%, jf. Tabel 1.



Figur 6: Værdien af gns. rejsetid (TT) og rejsetidsvariabilitet (TTV) i basisscenariet og case IV.

Opsummering af cases og sammenligning med nuværende metode

I alle fire cases ser vi, at omkostningerne forbundet med rejsetidsvariabilitet udgør en væsentlig del af de samlede rejsetidsomkostninger, og – mere vigtigt – at omkostningerne forbundet med den sparede TTV i forhold til basisscenariet udgør en relativt stor del af den samlede gevinst.

Tabel 2 viser rejsetidsomkostninger opgjort efter den nuværende metode, hvor de er baseret på værdien af fri køretid og gns. forsinkelse. En sammenligning med Tabel 1 viser, at den nuværende metode i disse simple scenarier undervurderer de totale omkostninger, men giver den samme rangordning. Betragter vi den relative gevinst i forhold til basisscenariet, ses at den stort set er ens for de to metoder for case I-III, mens der er noget større forskel mellem metoderne for case IV. Forskellen mellem metoderne (målt ud fra den relative gevinst) afhænger selvsagt af, hvor stor standardafvigelsen i rejsetiden er i forhold til den gennemsnitlige forsinkelse, men også af forholdet mellem den gennemsnitlige rejsetid og den fri køretid. Derudover afhænger den naturligvis direkte af den anvendte værdi for rejsetidsvariabilitet: En værdi der, skønt den er baseret på internationale erfaringer, må betragtes som et groft overslag, da disse internationale studier resulterer i højst forskellige værdier.

Tabel 1: Resume af case- og basisscenerier

	Andel dage med trængsel	Gns. Varighed af trængsel	Omkostninger fra gns. rejsetid	Omkostninger fra TTV	Samlede rejsetidsomkostninger
Scenarie		Antal kvarter	DKK/morgen /vej-km	DKK/morgen /vej-km	DKK/morgen /vej-km
Basisscenarie (case I-III)	88,9%	8,08	68.718	26.784	95.501
Case I. Peak spreading	67,8%	6,45	61.350	23.633	84.982
Case II. Rampedosering	82,6%	7,90	67.081	26.150	93.231
Case III. Fra 3 til 4 spor	14,5%	5,83	53.928	11.358	65.287
Basisscenarie (case IV)	52,4%	6,41	32.230	10.696	42.926
Case IV. Fra 2 til 3 spor	1,7%	6,20	28.021	2.813	30.834

Tabel 2: Rejsetidsomkostninger beregnet efter nuværende metode

	Omkostninger fra fri køretid	Omkostninger fra gns. forsinkelse	Samlede rejsetidsomkostninger
Scenarie	DKK/morgen/vej-km	DKK/morgen/vej-km	DKK/morgen/vej-km
Basisscenarie (case I-III)	51.958	24.797	76.756
Case I. Peak spreading	51.958	14.063	66.021
Case II. Rampedosering	51.958	22.376	74.334
Case III. Fra 3 til 4 spor	51.958	2.915	54.874
Basisscenarie (case IV)	27.890	6.423	34.313
Case IV. Fra 2 til 3 spor	27.890	195	28.085

Referencer

de Palma, A., Fosgerau, M. (2011): Dynamic Traffic Modeling, kap. 9 i *A Handbook of Transport Economics*. Edward Elgar Publishing, UK.

DTU Transport: Travel time variability - Definition and valuation, Rapport 1-2008.

1. Introduction

This report describes the development of a prototype model to predict travel time variability (TTV) on Danish motorways.

TTV measures the extent of unpredictability in travel times faced by travellers. Unpredictability can arise due to day-to-day fluctuations in traffic demand, traffic incidents affecting capacity or weather conditions. In this report, TTV is measured as the standard deviation of travel times over all typical weekdays in the period of analysis.

The model has been developed by DTU Transport for the Danish Ministry of Transport and the Danish Road Directorate, as one of the steps towards including travel time variability in cost-benefit analyses of transport projects in a theoretically satisfactory way.² The model takes as input a prediction of travel demand (stemming e.g. from the national traffic model, LTM) and simulates the travel time mean and variance.

The model is intended as a post-processing module to be applied after a traffic model has been used to predict travel demand on a road network. Its output (the level of TTV) does not feed back into the traffic model to account for the behavioural response to TTV. The model is thus an example of a 'Method 1' approach in the terminology of de Jong and Bliemer (2015), who recommend using this simple type of model in the short run. Such approaches have been implemented in the Netherlands, the UK, the US and Sweden. In the longer run, they recommend working towards implementing the prediction of TTV and the behavioural response to TTV into the traffic models.

The prediction model is based on a statistical model of the relationship between observed travel times and traffic flows. Both travel time and traffic flow are dynamic processes that evolve over the day and affect each other. Our aim has been a simple model which takes into account the dynamic relationship between the two and avoids the potential endogeneity issues related to this relationship. In this respect, our model constitutes a significant methodological improvement compared to the traditional speed-flow curves. The model and the estimation of its parameters is described in sections 3 and 4.

The prediction model enables us to predict mean travel time and TTV (measured as the standard deviation of travel time over different days) for a given traffic scenario, and to compute the travel costs associated with both, assuming a known monetary value of TTV. In section 5, we present four simple case scenarios to illustrate the use of the method. In the scenarios, we apply a reliability ratio of one, i.e. the value of one minute's standard deviation equals the value of one minute's mean travel time, cf. the recommendations in DTU Transport, 2008. In most scenarios, we find that the change in the costs of TTV makes a significant contribution to the overall change in travel costs.

² The current Danish practice for private transport distinguishes between free flow travel time and mean delay and values mean delay at 1.5 times the value of travel time, while for public transport it distinguishes between scheduled travel time and mean delay and values mean delay at 2.0 times the value of travel time. To the extent that the mean delays relative to free flow / scheduled travel time are related to TTV, this method does somewhat account for TTV. However, most likely the method does not account for the entire unpredictability of travel times. Hence, DTU Transport (2008) has recommended to implement a new practice, distinguishing instead between mean travel time and TTV (represented by the variance or standard deviation of travel times or a similar suitable measure which can be assigned an economic interpretation). This has so far not been implemented in practice because the currently used traffic models cannot handle TTV.

The prediction model is estimated on data from the Køge Bugt Motorway, a congested Danish motorway through the south-western suburban area of Copenhagen towards the city centre. In principle, the model is general enough to cover other roads, as it is highly simple. However, we recommend extending it by re-estimating its parameters on data from other roads when the necessary data become available. Moreover, we recommend further developing the model to take spillback effects into account. We discuss these points in more detail in section 6.

We remark that de Jong and Bliemer (2015) recommend predicting the standard deviation of travel time from predicted mean travel times or mean delays, and review several international studies who estimated such relationships. They prefer this approach over an approach where TTV is predicted based on travel demand, as in our model, because with the latter approach policies that do not affect demand will not affect TTV. They admit, however, that the approach of predicting TTV from mean delays has a similar drawback, since it does not allow policies to affect TTV if they do not affect mean delays. Another reason they reject predicting TTV from travel demand is that such a prediction model is likely to vary between different routes due to different locations of on- and off-ramps. However, in our opinion, there is no reason why the relationship between TTV and mean delays should not also vary between routes for the same reasons. In the end, it is necessary to make some assumptions about the distribution of travel times in our model, and these assumptions naturally impose a functional form on the relationship between TTV and mean delay, with parameters that are estimated from the data. So in this sense, the two approaches are similar. However, in applications, they will not be similar: The relationship between TTV and mean delay holds for the mean delays predicted by our dynamic model, which most likely differ from those predicted by the LTM, which stem from speed-flow curves that do not account for dynamic effects.

2. Data description

Our analysis uses traffic data from the Køge Bugt Motorway (E20 between Køge and Avedøre) in 2012-2013, direction towards Copenhagen, from the Road Directorate's Mastra and Hastrid systems. This choice of data is motivated in the section below. Section 2.2 describes other data sources used together with the Mastra-Hastrid data. Section 2.3 provides our definitions of relevant analysis variables. Section 2.4 concerns the selection of the analysis sample and section 2.5 provides some statistics.

2.1 Choice of main data source

Our goal was to have traffic data with:

- a) Lots of congestion and travel time variability (to enable an analysis).
- b) Data for a series of adjacent road segments (to ensure that we know traffic conditions on links upstream and downstream from the link we analyse, for use as potential instruments).³
- c) An adequate coverage to compute travel time variability: To compute variability at a given time of day t , we need observations from several days at time t . Moreover, to allow for identification of potential dynamic effects, we need to know travel conditions not just at time t , but also at specific time intervals prior to t (such as 15 and 30 minutes prior to t). In summary, we need several days with contiguous series of traffic information.

To keep our analysis as simple as possible, we had planned to analyse motorways only. Modelling dynamic speed-flow relations on other types of roads with turns and traffic lights is a complicated issue and still an area of ongoing research.

Initially, we had intended to use travel time data from the Road Directorate's TRIM system on the motorways in Trekantsområdet. The reason was that the TRIM system measures travel times at the road segment level, computed using cameras and licence plate recognition. However, since this part of the motorway network does not suffer from congestion in the same degree as the roads in the Copenhagen Area, we decided upon using other data sources. Note that this is not a question of prioritising between different parts of the country, but merely a question of securing data quality: To develop a model of travel time variability, it is of crucial importance that we use data where travel times display lots of variation over days.

Point c) above rules out most types of GPS data, as these do not provide sufficient coverage of a specific road link to compute reliable measures of travel time variability.⁴

Another option was to use Bluetooth data, where travel times are based on observations of Bluetooth devices at different locations along a road link. This is potentially a very strong data source: Though analysis of such data has its challenges with regards to e.g. vehicle identification and representativeness, they have many advantages: 1) They measure travel times at the level of the individual car (or Bluetooth device). 2) It is possible to obtain data with a good coverage of specific road links. 3) It is possible to measure travel times even in very congested conditions, where e.g. loop detectors cannot provide reliable measurements. We

³ As it turned out, we did not use this information in the analysis.

⁴ An exception is experimental data, where a GPS-tracked car fleet repeatedly travel the roads of interest. However, such data were not available.

consider using data from the project “Flaskehalse” (cf. Tetraplan 2013) which recorded Bluetooth travel times on the Køge Bugt and Helsingør motorways for an approximate 3 week period in 2013. We rejected this data partly because the survey period was rather short (as the aim was computing day-to-day variation), partly due to methodological reasons, as the “Flaskehalse” project considered only few and rather long road segments whereas we intended our analysis to be based on several adjacent road links, in line with the analysis in Fosgerau & Small (2012). However, since we ended up using a different methodological approach, it is possible that the data can be used as a supplementary check.

Finally, we decided on using loop detector/radar travel time data from the Road Directorate’s Hastrid system, which contain measurements of travel times for major Danish state-owned roads. Though loop detector data suffer from unreliable estimates of low speeds (<15km/h)⁵ and from observing travel times at the level of location points rather than at link level, they have the advantage of providing a good coverage in both time and space dimensions, for a long contiguous period. We combined the Hastrid data with travel flow data from the Road Directorate’s Mastra system, which contains traffic counts for a wide range of Danish roads. For both Hastrid and Mastra, the amount of data available and the level of detail vary greatly over different locations and different periods of analysis, which naturally affected our choice of locations for analysis.

We decided on two of the “arterial” motorways leading to/from Copenhagen, as these are among the most congested in Denmark and (as a consequence) quite well covered data-wise. As recommended by the Road Directorate, we preferred the analysis period to be as recent as possible, since newer data is supposedly of better quality.

As part of the project, we generated datasets for the Køge Bugt Motorway and the Helsingør Motorway (direction towards Copenhagen) for the period 2012-2013. As both database setup and analysis took longer time than anticipated, we were however forced to limit our analysis to the Køge Bugt Motorway, due to time limitations.

2.2 Additional data sources used in analysis

To provide information about road works during the analysis period, we use data from Trafikman, a real-time traffic information system operated by the Road Directorate. Trafikman provides traffic information to car drivers via the Road Directorate’s webpage, traffic radio stations, and smartphone app’s. Each data record consists of a time stamp, an (approximate) location and an incident description from which we can infer incident type (planned road works, unplanned maintenance work, accident, queue, dropped items blocking the road, etc.). Incidents are indexed, such that it is possible to identify all records relating to an incident. There is a record for each update regarding the incident in question, including records of when incidents have ended.

Our definition of road work is broad and encompasses both planned works and unplanned maintenance work, since we are interested in all types of work that may temporarily block one or more of the lanes. We defined road work as an incident:

- with type ‘ROV’ (=road work)⁶, or

⁵ We were advised to discard observations with speeds below 15km/h, though a specific technical threshold is unknown: The problems with unreliable speed may also apply to some degree to measured speeds above 15 km/h.

⁶ The incident type ‘ROV’ was implemented during 2013, so all road work incidents in 2012 were of type ‘DRD’.

- with type 'DRD' (=Danish Road Directorate), where the description contained one of the Danish keywords:
 - "vejarbejde", "byggearbejde", "slag hul" or
 - "vejbelægning i dårlig stand" in combination with "vejhjælp er tilkaldt" or "vejhjælp er på vej".

Finally, we have available weather data (temperature, precipitation, wind speed, sight, and snow depth) from four weather stations located at Roskilde Airport, Roskilde, Central Copenhagen and Copenhagen Airport. These data stem from the Danish Meteorological Institute.

2.3 Definition of analysis variables

From the Hastrid system, the Road Directorate provided travel times at the road segment level for a pre-defined segmentation of both motorways. This yielded 8 road segments for the Køge Bugt Motorway, cf. *Figure 1*. We refer to these road segments as links. The links varied between 2 and 6 kilometres in length (roughly). The travel time data were available at the 1-minute level. However, since traffic flow data were available only at the 15-minute level, we aggregated travel time data to this level as well.



Figure 1: The eight road segments (links) on the Køge Bugt Motorway, separated by red dots. We consider the direction towards Copenhagen, i.e. from South-West towards North-East.

We extracted traffic count data from the Mastra system for all counting locations along the two motorways and used them to compute traffic flow:

1. First, we computed traffic flows (at the location point level) as #vehicles per lane per minute.
2. Then we converted this flow into #passenger-car-equivalents (pce) per lane per minute, by using appropriate conversion factors for vehicles of length 580-1250 metres and 1250- metres and vehicle length shares obtained from the nearest available location.⁷ The locations where vehicle length

⁷ The number of vehicles of length 580-1250 metres was multiplied by 1.5 and the number of vehicles of length >1250 metres by 2.0., cf. The Danish Road Directorate (2010a).

information is registered are listed in Table 3. Since vehicle length information was available only at very few locations, this nearest location may be several kilometres away, so the conversion should be seen as an approximation rather than an exact calculation.⁸ Moreover, the vehicle length information is often only available at the 1-hour level. Despite these circumstances, we believe that using an approximated #pce per lane per minute is still better than using the exact #vehicles per lane per minute.

3. Finally, we aggregated flows from the location point level to the link level by simply averaging over all available counting locations within each link. This procedure has the advantage of providing as much data as possible: We obtain data for a link even if data is missing from some of the locations along the link. Moreover, the method is consistent with the way the Road Directorate computes the link travel times, as these are computed as averages over all non-missing speed measurements in the given time interval (The Danish Road Directorate, 2010b). The drawback is that the exact definition of link flow differs over time, because i) some counting locations may be temporarily out of order in some periods, and ii) not all counting locations register traffic every day – some only register, e.g., three days each week or every second week. This is particularly problematic for link 7 of the Køge Bugt Motorway, which is why we chose to focus on the remaining links. To mitigate the potential problems, we left out counting locations with very few observations in the aggregation procedure, and defined indicator variables with information about flow definition such that we could keep track of the different definitions in our analysis (see next section about data selection). For future analyses, we strongly recommend using both travel times and counts at the location level, rather than aggregating to link levels.

Table 1: The eight links on the Køge Bugt Motorway (direction Køge-Copenhagen)

Link no	Description	Length (km)
1	<32> Køge/Ølby - <31> Solrød S/Roskilde	3.88
2	<31> Solrød S/Roskilde - <30> Solrød	3.19
3	<30> Solrød - <29> Greve S	5.44
4	<29> Greve S - <27> Hundige	3.59
5	<27> Hundige - X Ishøj	1.68
6	X Ishøj - <25> Vallensbæk	4.30
7	<25> Vallensbæk - X Avedøre	3.16
8	X Avedøre - <22> Gl. Køge Landevej	1.89

⁸ Note that for links 6 and 8, there are in principle closer registrations available, from September 2013 onwards. However, this is not relevant for link 6, as the analysis sample for other reasons was restricted to be before September 3rd 2013 for this link (cf section 2.4). For link 8, we decided to use the locations in Table 3 throughout the analysis period, for the sake of consistency. It is relevant to note however that the actual vehicle length shares on link 8 may deviate somewhat from those at the locations in Table 3.

Table 2: The Mastra counting locations used when computing link flows

Link no	Mastra counting location
1	0 10-0 34/ 145 -
2	0 10-0 29/ 500 -
2	0 10-0 30/ 400 -
3	0 10-0 26/ 300 -
3	0 10-0 28/ 200 -
3	0 10-0 28/ 750 -
4	0 10-0 20/ 620 -
4	0 10-0 21/ 400 - *
4	0 10-0 22/ 540 -
5	0 10-0 18/ 600 -
5	0 10-0 20/ 620 -
6	0 10-0 15/ 700 - *
6	0 10-0 16/ 0 -
6	0 10-0 16/ 500 - *
6	0 10-0 17/ 1 -
6	0 10-0 17/ 60 - *
6	0 10-0 18/ 600 -
7	0 10-0 11/ 200 - **
7	0 10-0 12/ 0 - **
7	0 10-0 12/ 500 - **
7	0 10-0 13/ 540 - **
7	0 10-0 13/ 541 - **
7	0 10-0 14/ 590 - **
7	0 10-0 14/ 591 - **
8	0 3-0 54/ 700 +
8	0 10-0 11/ 140 -
8	0 10-0 11/ 200 -

* Not used in current analysis, cf. section 2.4.

** Not used in current analysis, cf. section 2.3.

Table 3: The Mastra counting locations used for vehicle length information

Link no	Mastra vehicle length registration location, in order of preference *
1, 2, 3	0 10-0 26/ 300 – 0 10-0 22/ 540 – 0 10-0 20/ 620 –
4, 5, 6, 8	0 10-0 20/ 620 – 0 10-0 22/ 540 – 0 10-0 26/ 300 –

* Order of preference A-B-C means that A is used if available for the entire AM-period of a given day, otherwise B is used if available for the entire AM-period, otherwise C is used if available for the entire AM-period, otherwise the day is left out of analysis.

2.4 Sample selection

As mentioned, our analysis is limited to the Køge Bugt Motorway. We focused on modelling the AM-period (4:45 AM – 12 noon) for the direction towards Copenhagen, which has a distinct morning peak with lots of congestion.

We use observations between 4:45 AM and noon from weekdays (Mon-Fri) that are either characterised as “Typical weekdays” or “Special days” (but not holidays). We exclude observations from time intervals where we do not observe vehicle length shares at any location on the Køge Bugt Motorway. We also exclude observations from periods with road works, using the traffic information data from Trafikman, since we want to know the number of lanes available.⁹

We exclude observations where travel time (per kilometre) exceeds 4 minutes, corresponding to speeds below 15 km/hour, as the Hastrid data are not reliable for such low speeds. We also exclude observations where the traffic flow exceeds 40 pce per lane per minute, as these appear to be outliers.

The Hastrid data contains many cases where the measured travel time is constant over several contiguous 15-minute intervals. This is particularly frequent in links 7 and 8. We are not aware of the reason for this, but due to the risk that it is caused by malfunctioning equipment we have chosen to exclude all observations where travel time is constant over more than two contiguous intervals and all observations from days with more than 4 constant-travel-time-spells each lasting at least two intervals.

To ensure data quality, we also impose some restrictions on the counting locations used to compute link flows (see Table 4 below). Moreover, we excluded observations before January 21st 2012 on links 1, 2 and 3, because the share of long vehicles differed systematically from the remaining period. Finally, the data for links 4, 5, and 6 reveal a systematic change in the pattern of travel times around the start of September 2013. This may be related to the replacement of equipment at a couple of measurement locations, but we cannot be sure of the cause and so decided to leave out observations after September 2nd 2013 for these links. For link 8, travel times in the period Jan 1st - Nov 8th 2012 are systematically wrong (according to the Road Directorate), and so were excluded from our analysis. Moreover, travel times on link 8 in the period Nov 9th – Dec 16th 2012 are systematically different from the remaining period while the observed flows do not appear to change. We excluded observations from the affected period, in case the pattern is caused by malfunctioning equipment.

Table 4: Restrictions imposed in sample selection

Link	Link-specific sample restrictions applied in analysis
1	Observations before Jan 21 st 2012 are excluded (share of long vehicles).
2	Observations before Jan 21 st 2012 are excluded (share of long vehicles).
3	Count from either “0 10-0 26/ 300 -“ or “0 10-0 28/ 200 -“ should be non-missing, as these are measured after the merging at entry ramp <30>. Observations before Jan 21 st 2012 are excluded (share of long vehicles).
4	Counts from both “0 10-0 20/ 620 -“ and “0 10-0 22/ 540 -“ should be non-missing, to ensure we observe flow both before and after entry ramp <28>. Observations after September 2 nd 2013 are excluded (unexplained systematic change).
5	Counts from both “0 10-0 18/ 600 -“ and “0 10-0 20/ 620 -“ should be non-missing, to

⁹ For links 3-8, that are not directly upstream from a very congested link, we exclude only observations with road works on the current link. For links 1 and 2, which are very congested, we can observe that traffic conditions are affected by conditions downstream. So for link 1 we exclude observations with road works on either link 1, 2 or 3, and for link 2 we exclude observations with road works on either link 2 or 3.

	ensure we observe flow both before and after entry ramp <27>. Note that we do not observe counts from the two lanes merging with O4, so we have to assume that the flow (per lane) is the same for these two lanes as for the three lanes remaining on E20. Observations after September 2 nd 2013 are excluded (unexplained systematic change).
6	Counts “0 10-0 16/ 0 –” and either “0 10-0 17/ 1 -” or “0 10-0 18/ 600 -” should be non-missing, to ensure observe flow both before and after entry ramp <26>. Observations after September 2 nd 2013 are excluded (unexplained systematic change).
7	<link not used in analysis>
8	Counts from both “0 3-0 54/ 700 +” and “0 10-0 11/ 140 -” should be non-missing, to ensure we observe flow both before and after the merging with E47. Observations in the period Jan 1 st 2012 - Nov 8 th 2012 are excluded: According to the Road Directorate the travel times from Hastrid from this period are likely to be systematically wrong. Observations in the period Nov 9 th – Dec 16 th 2012 are excluded, since this period involved unusually large fluctuations in the registered travel times.

2.5 Descriptive statistics

Figure 2 - Figure 5 show means and standard deviations of travel time and travel flow, as a function of time of day (mean and standard deviation is computed over different days). For links 1, 2, 3 and 8, we see a distinct morning peak in both mean flow and flow variability, together with a distinct peak in mean travel time and travel time variability. For links 4, 5 and 6, the pattern for mean travel time is much less pronounced: The mean travel time during the morning peak is only slightly higher than outside the peak. Travel time variability does increase somewhat during 8AM-9AM, but not much compared to the other links.

Figure 6 shows speed-flow plots for the seven road links. It is interesting to notice that the plots have the characteristic backward bending form for links 1, 2, 3 and 8, where mean travel time is highly affected by the morning peak congestion, while the plots for links 4, 5 and 6 miss the congested lower “branch” of points.

To understand the dynamics of congestion, it is relevant to study plots as *Figure 7-Figure 13*, which show travel time and flow profiles over some given mornings in the sample. The Figures reveal substantial variation between days, but it is possible to make out a simple general pattern: On days without congestion, travel time seems unrelated to the flow. On days with congestion, flow increases from a low value in the early morning to a high value at the beginning of the peak. Before the peak starts, travel time is stable at a low value and seems to be unrelated to the flow. However, once flow reaches a certain level, travel time increases sharply and remains high (and highly variable) during some time, until it decreases slowly to a low and stable midday level. This stylised pattern is depicted in *Figure 14* below and used to motivate the model we develop in section 3.

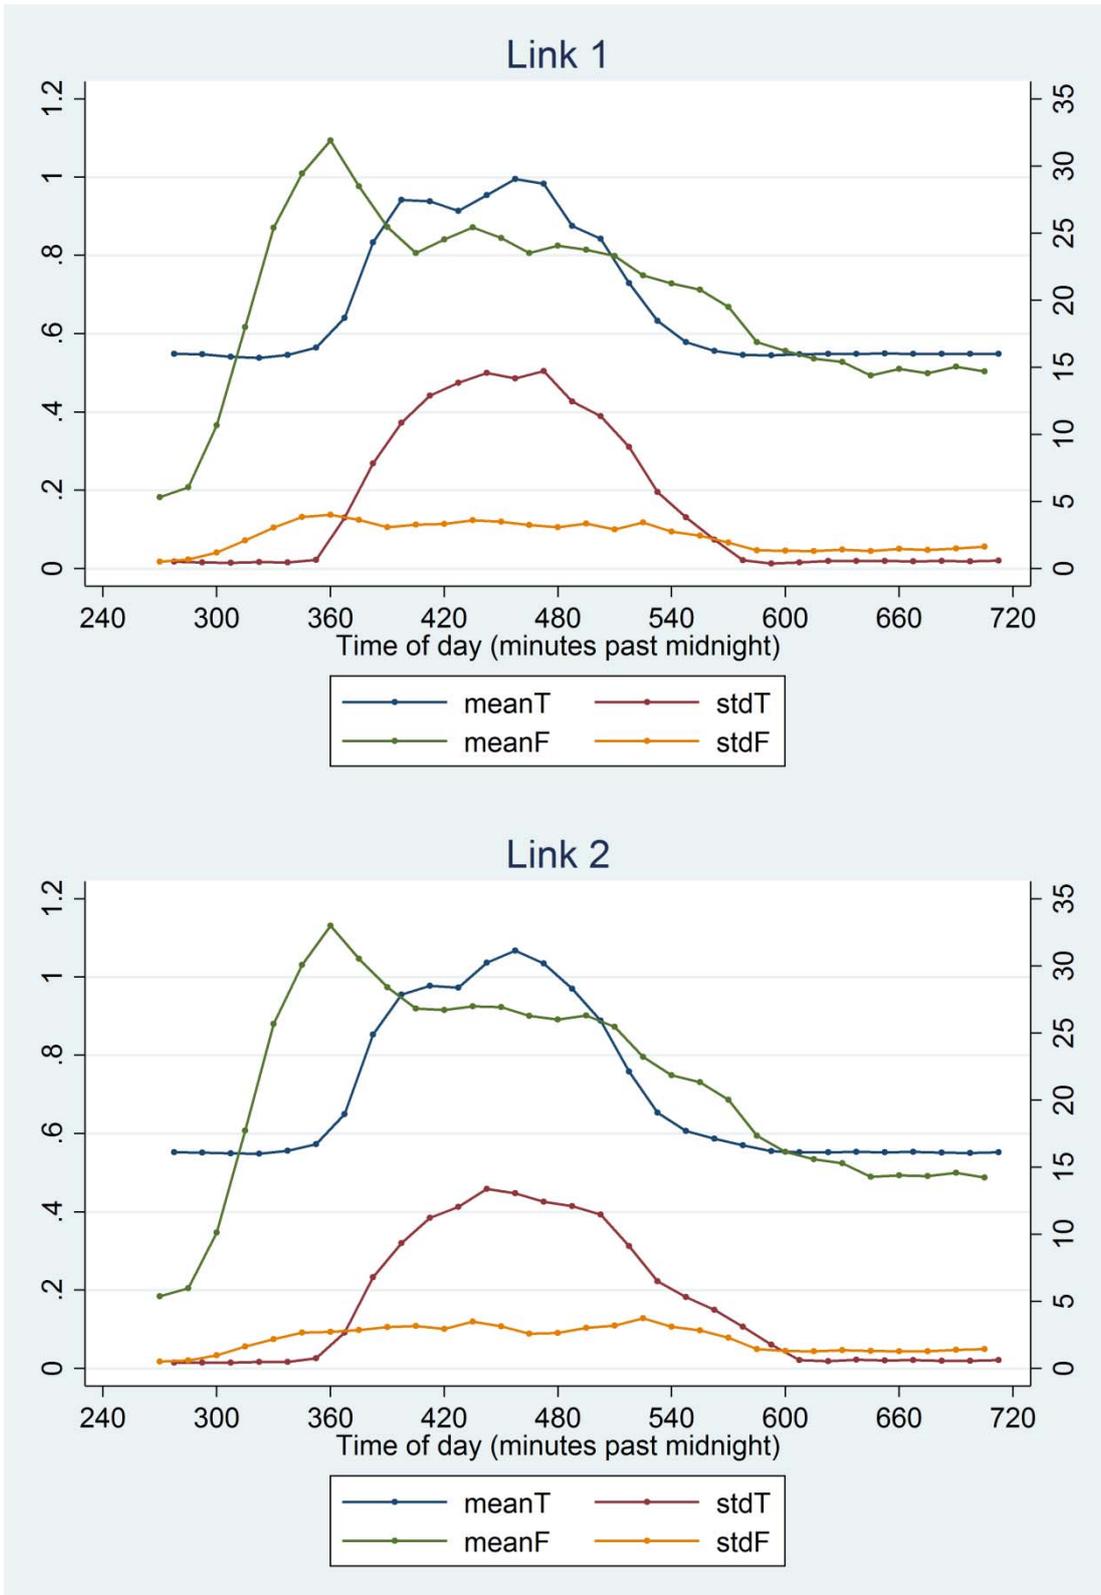


Figure 2: Mean and standard deviation of travel time T (in min/km) and flow F (in pce/lane/minute)

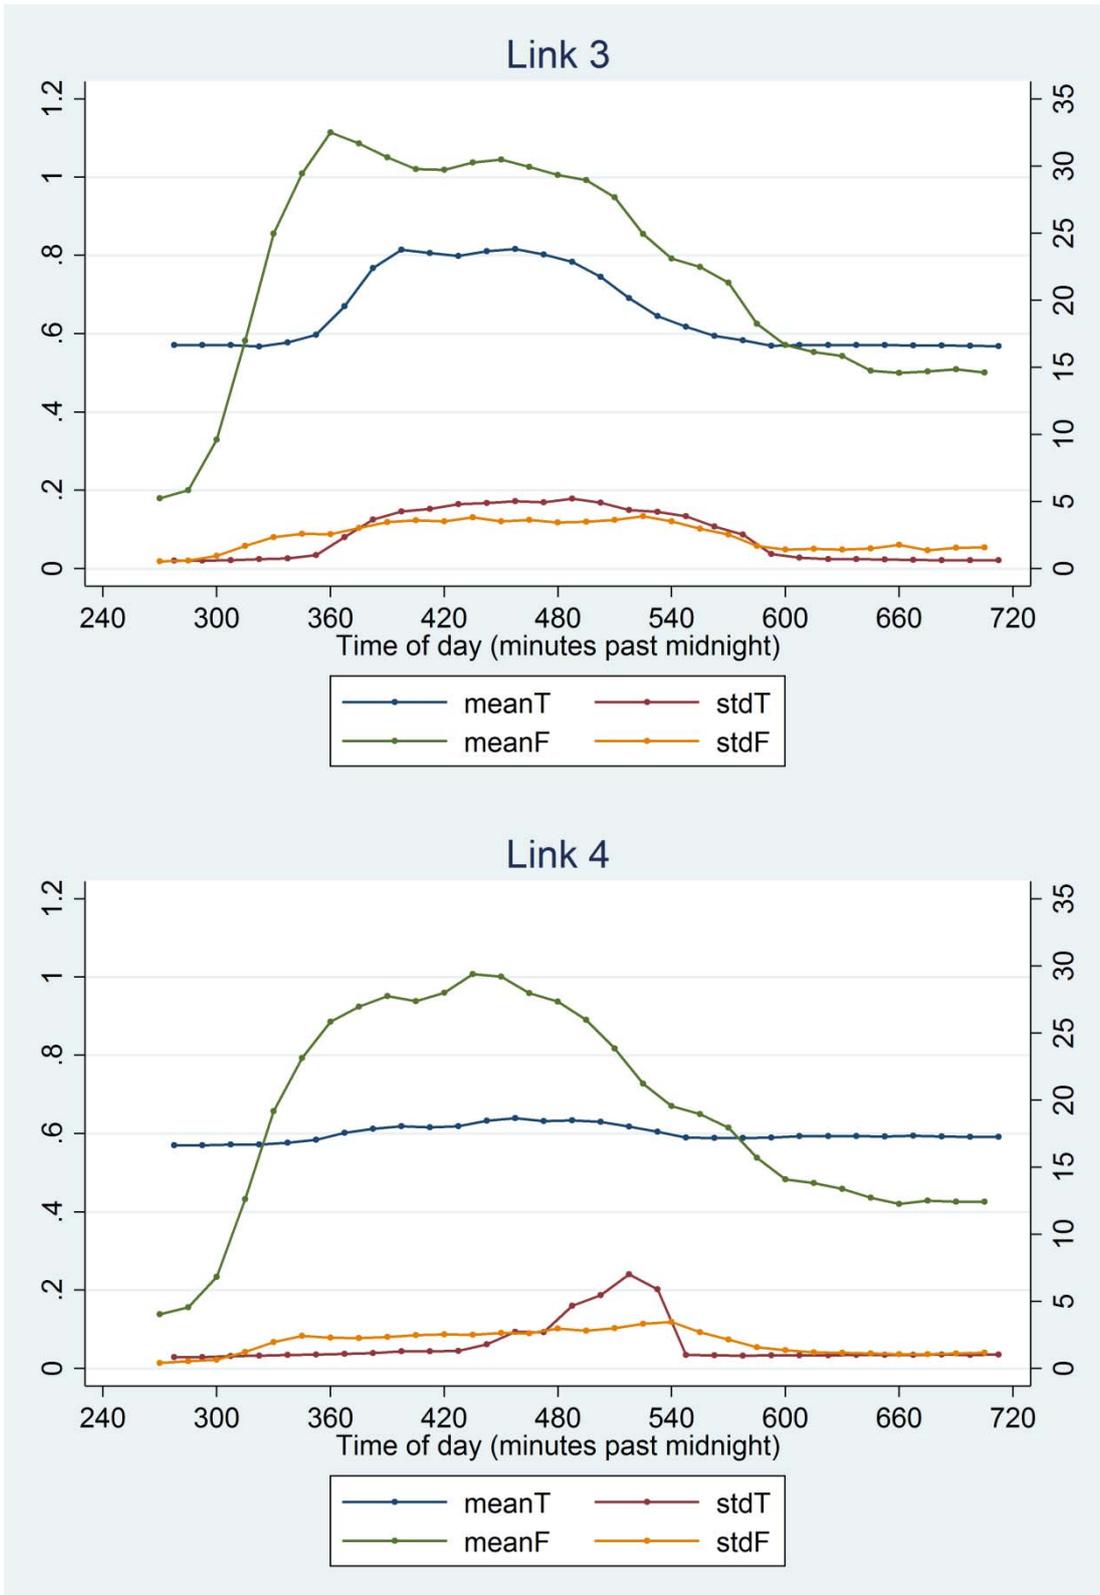


Figure 3: Mean and standard deviation of travel time T (in min/km) and flow F (in pce/lane/minute)

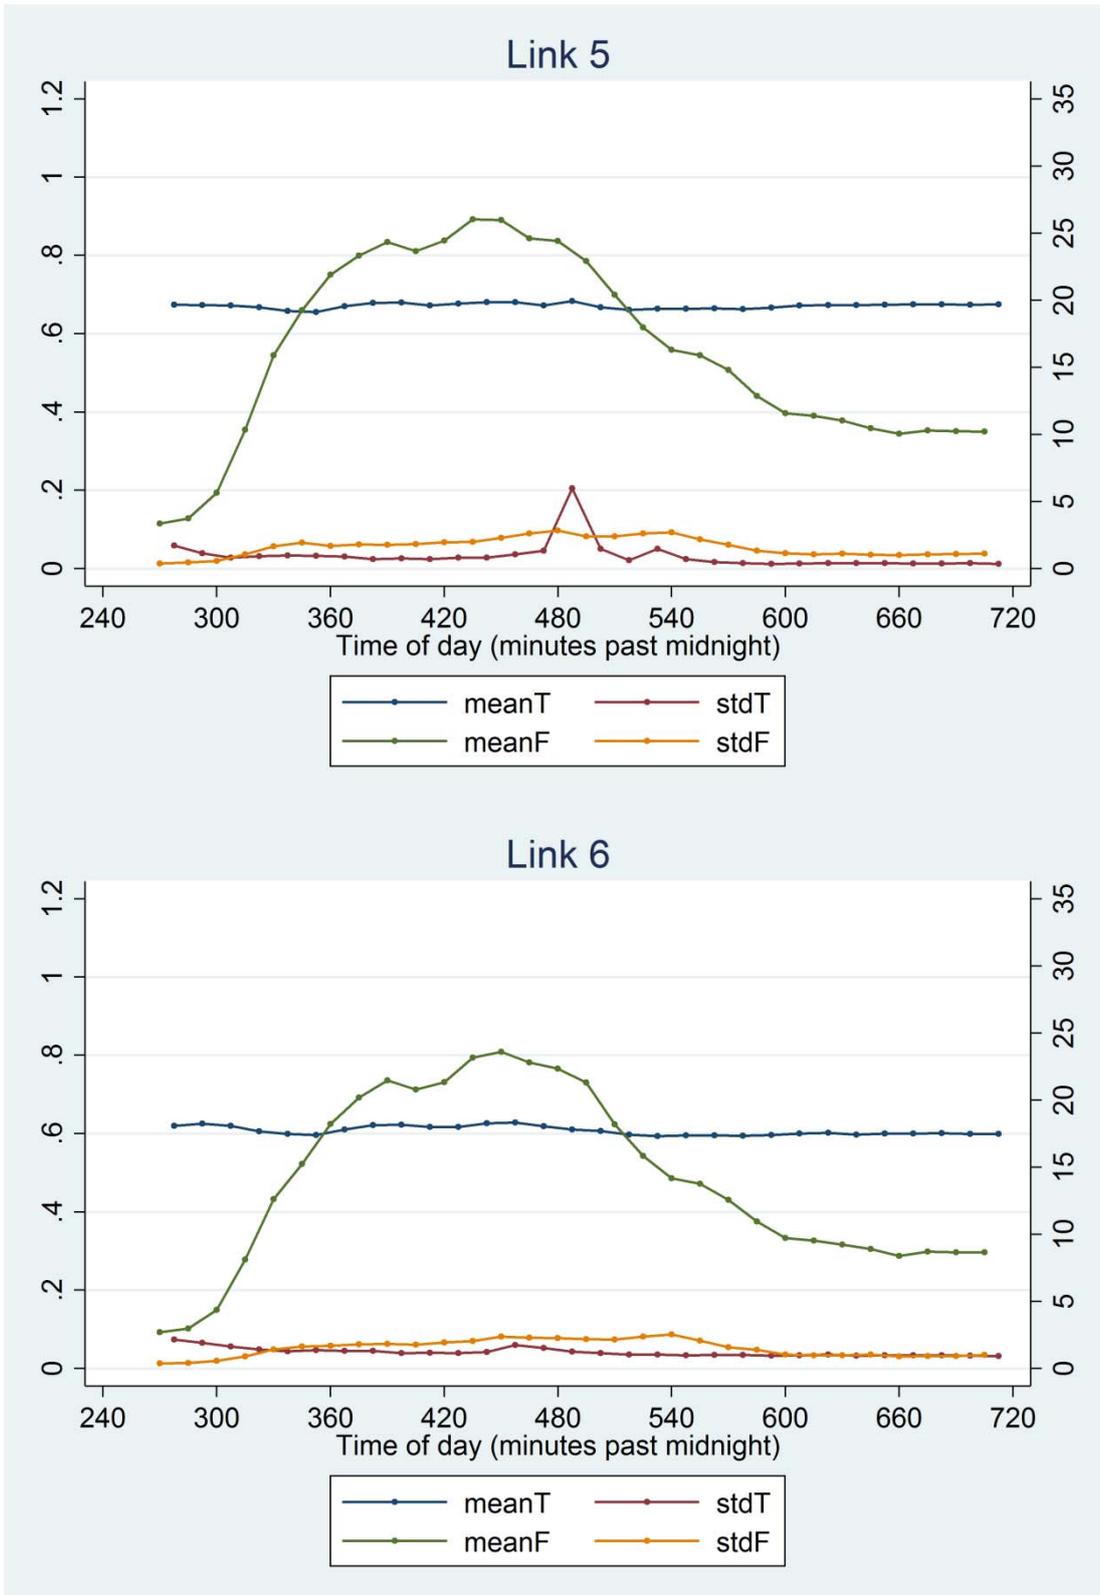


Figure 4: Mean and standard deviation of travel time T (in min/km) and flow F (in pce/lane/minute)

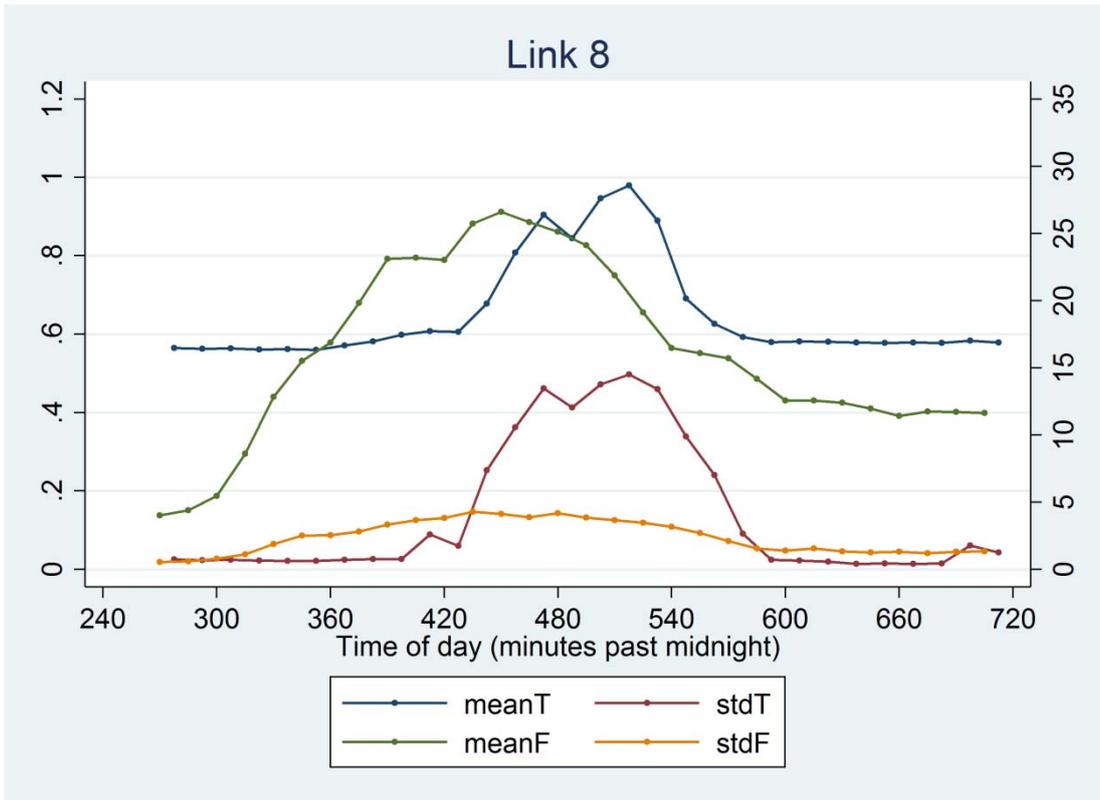


Figure 5: Mean and standard deviation of travel time T (in min/km) and flow F (in pce/lane/minute)

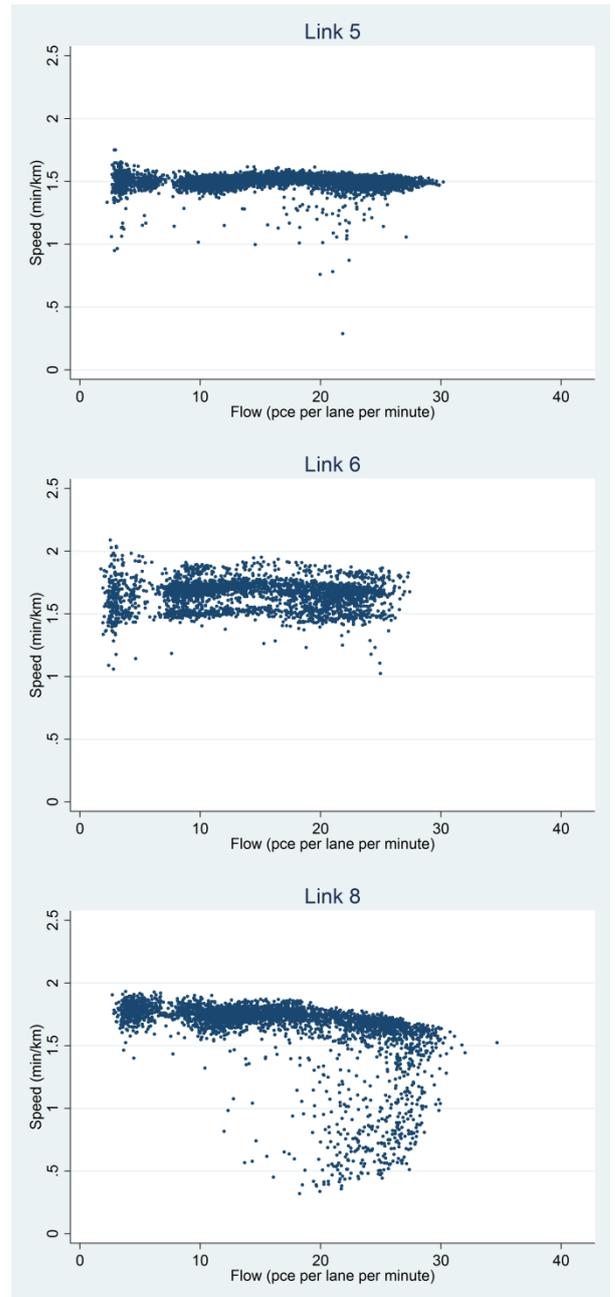
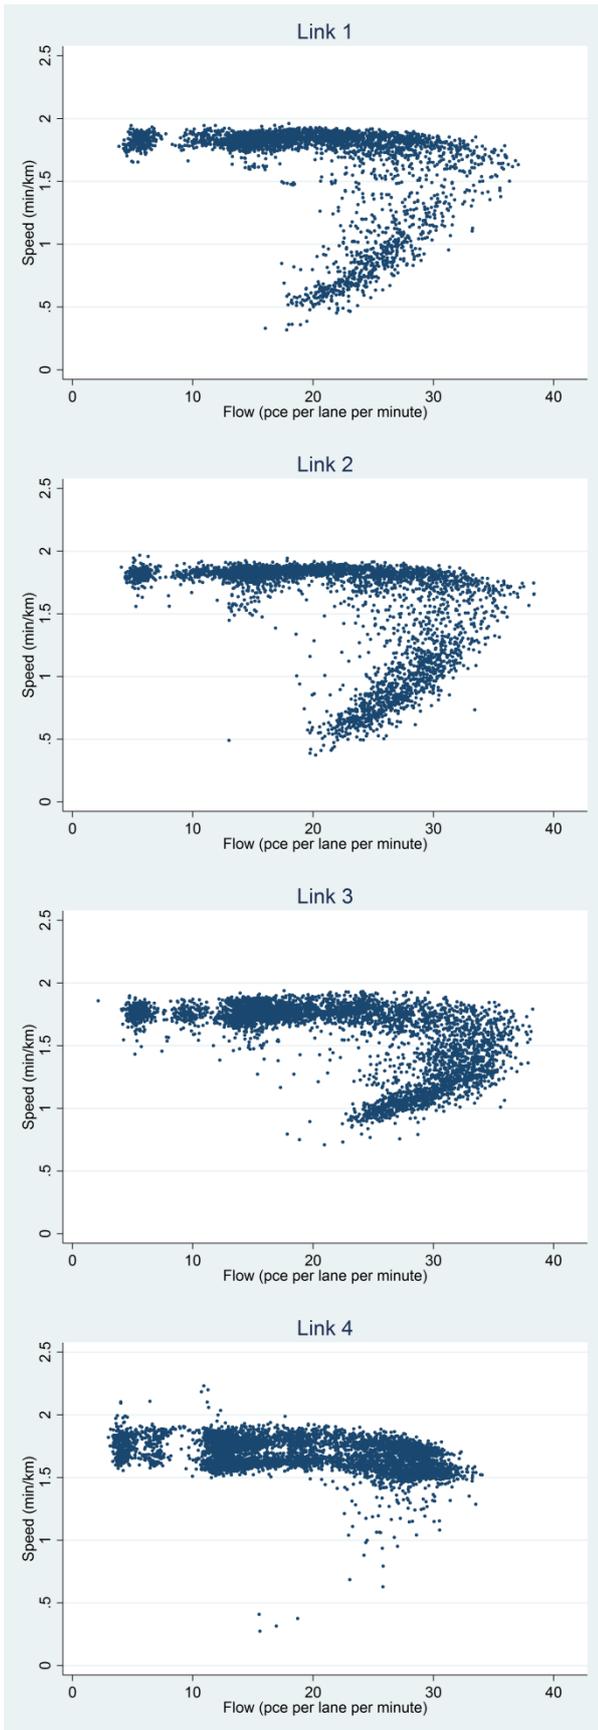
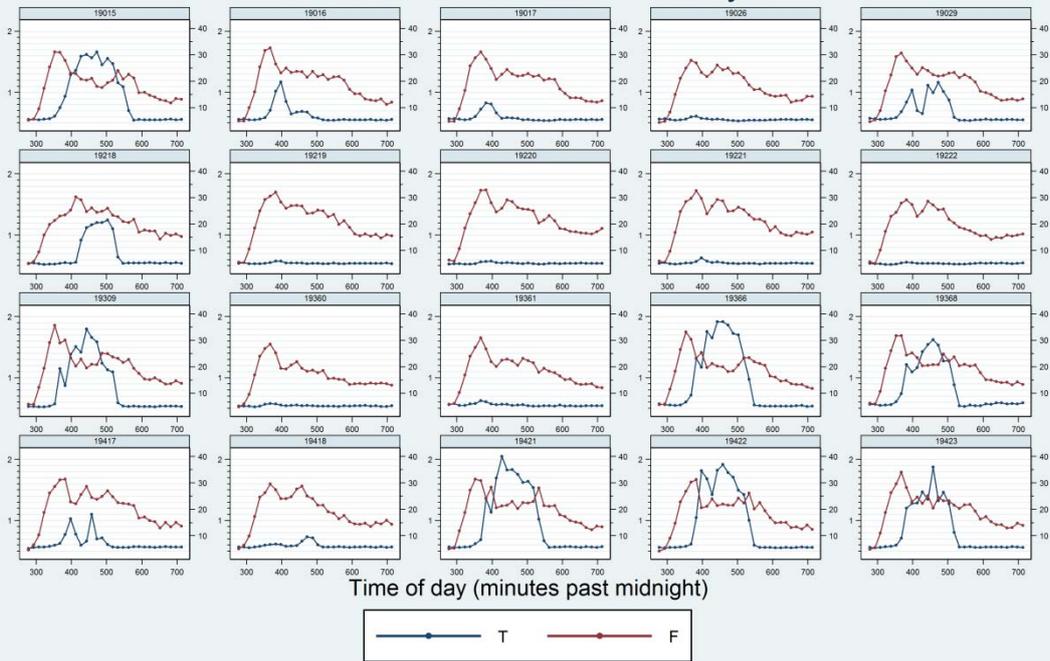


Figure 6: Speed-flow plots for the links

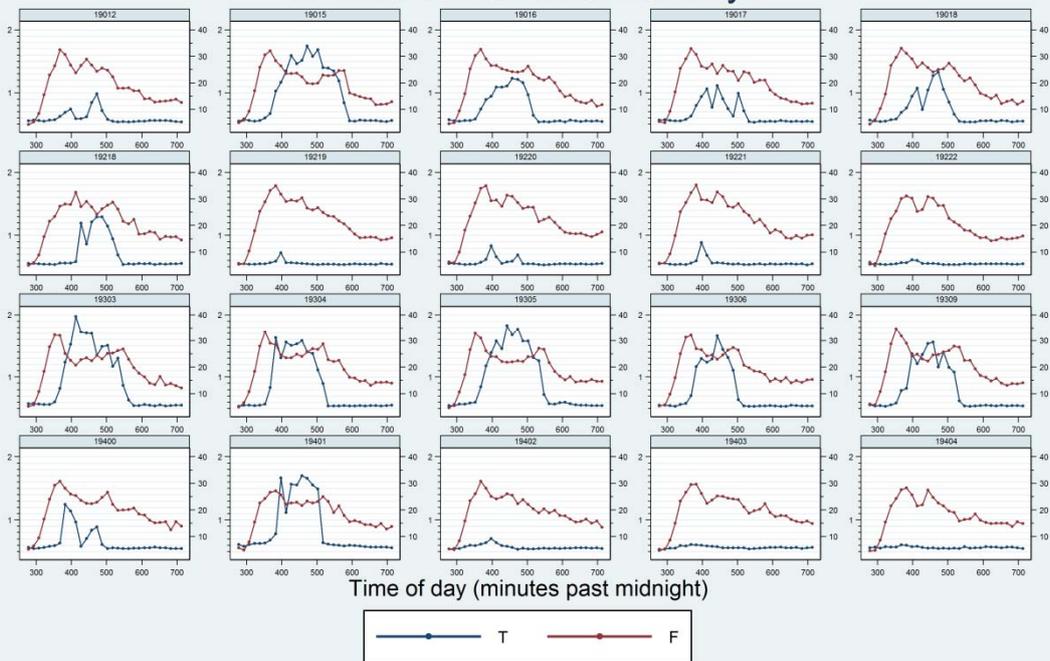
Link 1: Some different days



Graphs by Date

Figure 7: Examples of daily patterns of travel time and flow (link 1)

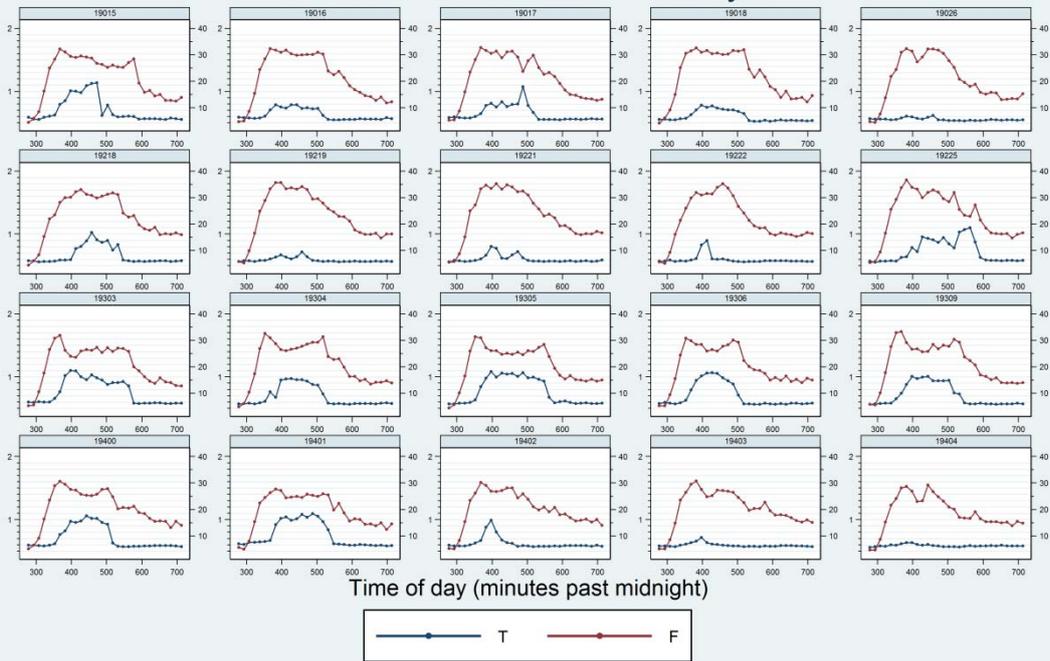
Link 2: Some different days



Graphs by Date

Figure 8: Examples of daily patterns of travel time and flow (link 2)

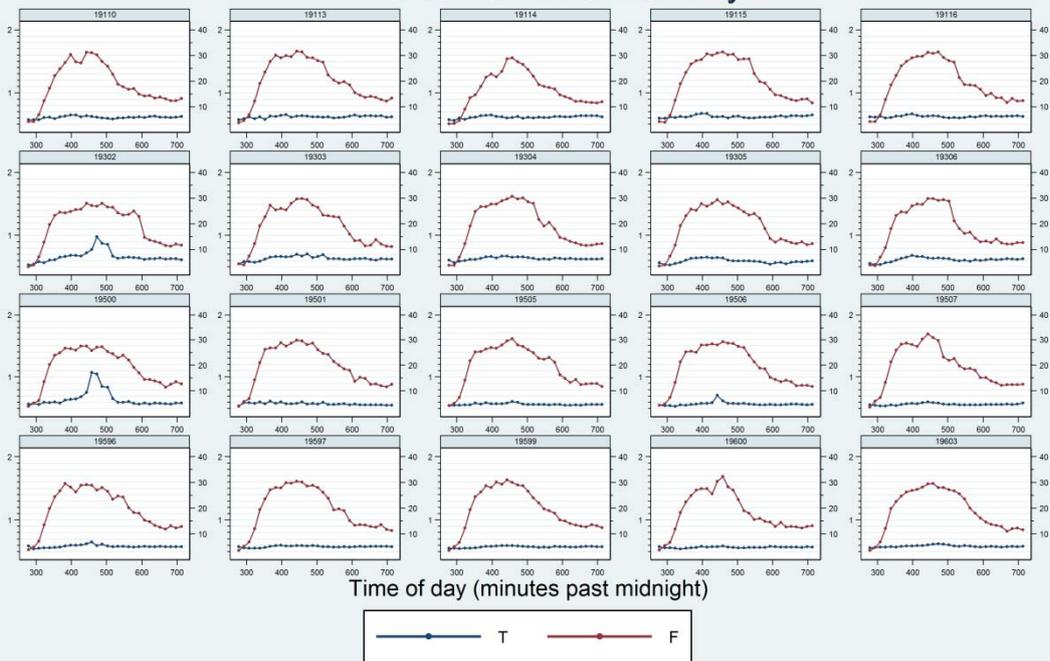
Link 3: Some different days



Graphs by Date

Figure 9: Examples of daily patterns of travel time and flow (link 3)

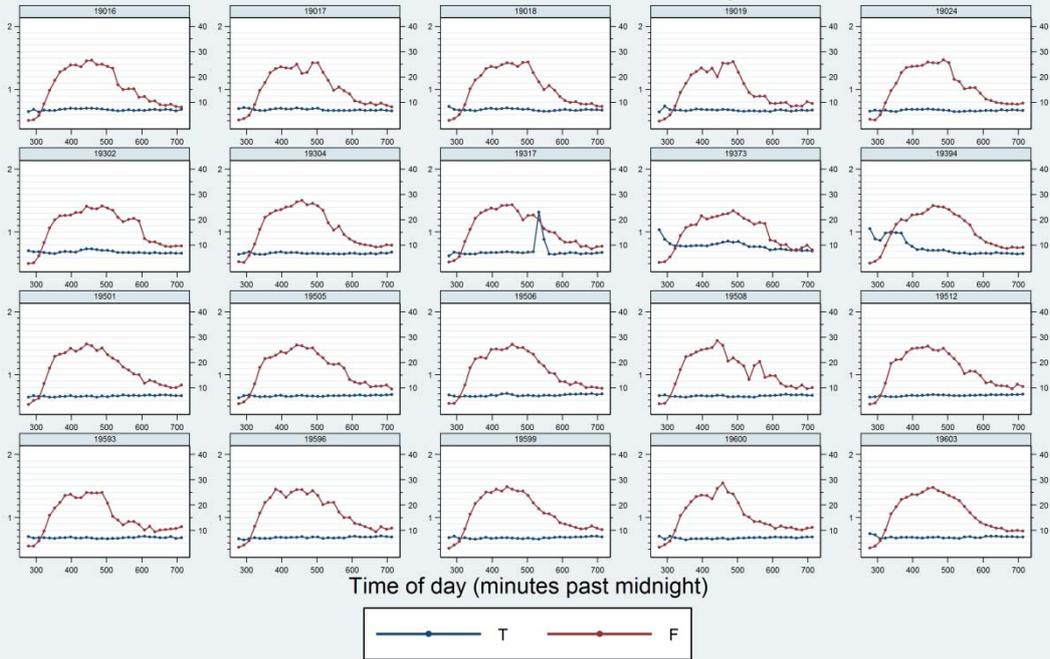
Link 4: Some different days



Graphs by Date

Figure 10: Examples of daily patterns of travel time and flow (link 4)

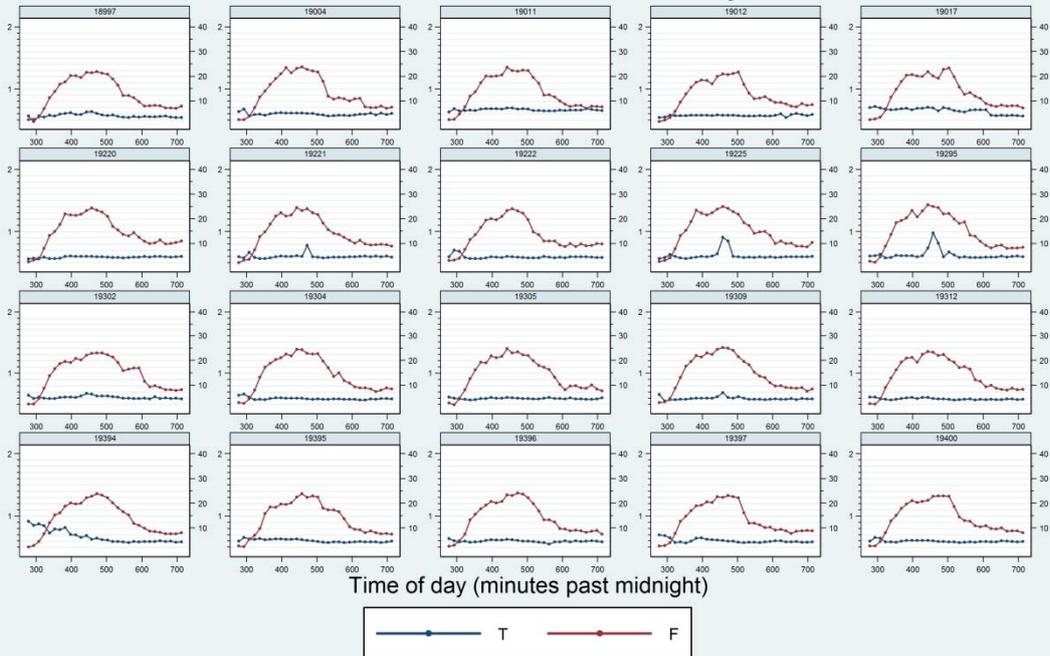
Link 5: Some different days



Graphs by Date

Figure 11: Examples of daily patterns of travel time and flow (link 5)

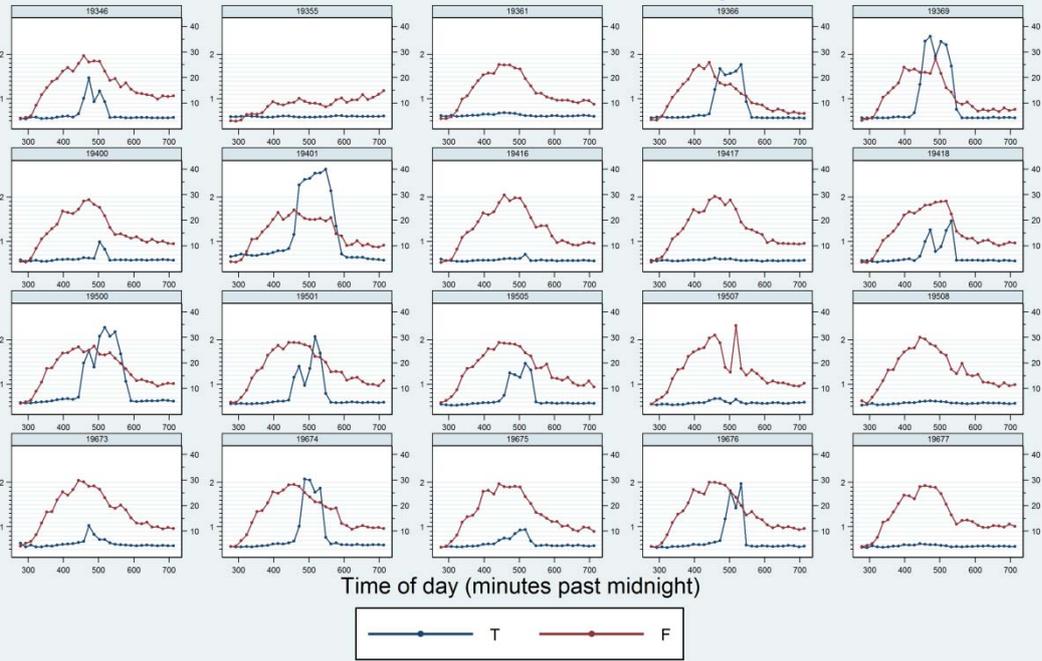
Link 6: Some different days



Graphs by Date

Figure 12: Examples of daily patterns of travel time and flow (link 6)

Link 8: Some different days



Graphs by Date

Figure 13: Examples of daily patterns of travel time and flow (link 8)

3. Model

The following model development is motivated by the observed patterns of travel time and flow over the day (described in the preceding section).

Assume two traffic states: Congested (c) and uncongested (u). Each morning traffic starts out in the uncongested state. It either stays in the uncongested state for the entire AM-period (from early morning until noon), or switches to the congested state sometime during the period, remains there for some time, and then switches back again. We assume that it switches to the congested state and back again *at most once*. We refer to the uncongested→congested transition as *breakdown* and to the congested→uncongested transition as *recovery*.

We further assume that travel time follows one distribution with cumulative distribution function (CDF) Φ_u in the uncongested state and another distribution with CDF Φ_c in the congested state.

Our model consists of three separate parts: A model predicting if and when breakdown occurs, a model predicting when recovery occurs – given breakdown time, and a model predicting means and variances of travel time depending on state.

Throughout, we index the 15-minute time intervals by their *end* time t , measured in minutes past midnight. The time interval [5:00AM; 5:15AM] is thus indexed by 315, and the last interval [11:45AM; 12 noon] is indexed by 720. In the model, T_t is the travel time per km, and F_t is the flow per lane per minute, in time interval t .

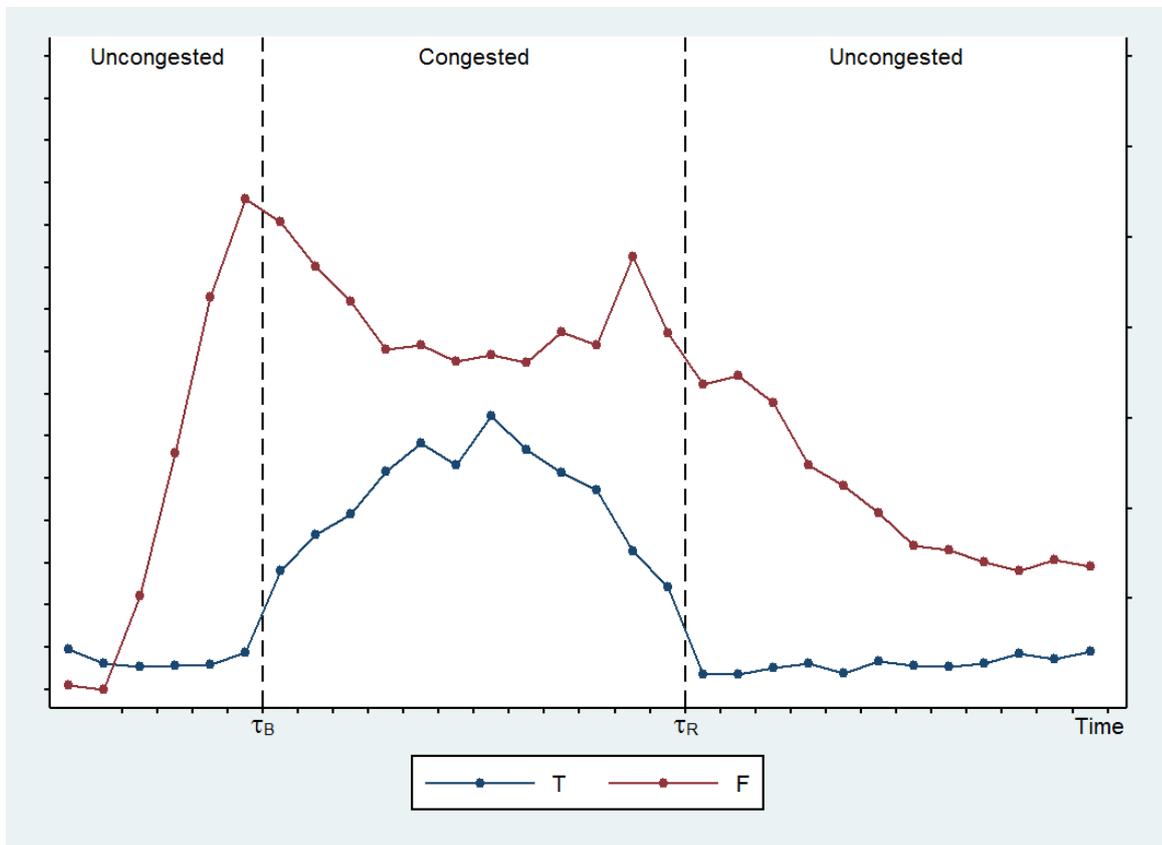


Figure 14: Illustration of traffic states. τ_B indicates breakdown time and τ_R recovery time.

3.1 Central concepts and assumptions

The following concepts are central to our analysis:

- The traffic demand D : The number of vehicles (per lane per minute) attempting to travel at the link.
- The exit flow: The number of vehicles (per lane per minute) leaving the link.
- The link flow variable F used in our analysis: A measure of the number of vehicles (per lane per minute) travelling the link within a 15-minute period, obtained as an average of point observations along the link (cf. section 2.3).

Consider a short link (meaning free flow travel time is well below 15 minutes). In uncongested circumstances the average exit flow in a 15-minute period will approximately be equal to the average demand in the period. In congested circumstances, average exit flow may be less than the average demand, since not all vehicles attempting to travel are able to get through: Above a certain point, increasing the number of vehicles on the link will result in reduced speeds, which will cause exit flow to decrease, since at lower speed fewer vehicles leave the link within a given time interval. This interdependency between speed and exit flow complicates statistical analyses of the relation between the two: When modelling speed as a function of exit flow one must take account of exit flow being endogenous, otherwise results will be biased.

As the link flow variable F used in our analysis is an average of point observations along the link, it is neither equal to the demand or the exit flow. Its interpretation is similar to the exit flow, however, since it is a measure of the number of vehicles actually passing through the link. Hence we expect it to behave similarly, i.e. we expect that under uncongested circumstances F will be equal to the average demand within the 15-minute period, and under congested circumstances F will be an endogenous predictor of speed.

More specifically, we make the following assumption in our analysis:

Assumption 1: Before breakdown, F is exogenous (to speed) and equal to average demand in the 15-minute interval. This seems a fair assumption, as travel times in uncongested periods (early morning and during middays) with low mean travel time and low travel time variability, appear to be unrelated to observed traffic flow. The assumption implies that we can use F to predict traffic states without worrying about endogeneity problems, and that we can use the estimated prediction model to explain how demand (from a traffic model) affects the breakdown probability.

3.2 Breakdown Model

We use a simple duration model to model the duration until breakdown. The dependent variable is the (discrete) breakdown time τ_B , which we define to be the start time of the 15-minute interval in which congestion starts (or, equivalently, the end time of the last 15-minute time interval before congestion sets in).

We model the probability that breakdown occurs at the end of time interval $]t - 15; t]$, conditional on not occurring before, i.e.:

$$P_B(t) = P(\tau_B = t | \tau_B \geq t) = f(x_t) \tag{1}$$

Here, x_t is a vector of explanatory variables related to the time interval $]t - 15; t]$. Note that this probability does not depend on clock-time or duration of the uncongested state. This is a deliberate choice: Though the breakdown patterns we observe in the data show a close relationship with clock-time, this effect works through the demand (flow) pattern and should be modelled as such to be of use in prediction.

We assume that $P_B(t)$ depends on traffic demand in time intervals prior to time t . Following Assumption 1, we estimate the effect of demand using observed flow F_t instead of demand. Hence the vector x_t includes functions of F_t .

3.3 Recovery Model

We also use a duration model to model the duration until recovery, after breakdown has occurred. The dependent variable is the (discrete) recovery time τ_R , which we define to be the end time of the last 15-minute time interval before recovery.

We model the probability that recovery occurs at the end of time interval $]t - 15; t]$, conditional on not occurring before, and conditional on $\tau_B = t_B \leq t - 30$, is:

$$P_R(t|\tau_B = t_B) = P(\tau_R = t | \tau_R \geq t, \tau_B = t_B) = g(z_t(t_B)) \quad (2)$$

Here, $z_t(t_B)$ is a vector of explanatory variables related to time interval $]t - 15; t]$ that may depend on breakdown time. To provide a theoretical basis for the choice of variables in $z_t(t_B)$, we consider the so-called bottleneck model from the literature of transport economics. This model describes the in- and outflow for a single bottleneck (a location with limited capacity) in an equilibrium situation where all travellers have identical preferences and no travellers can lower their generalised travel costs by shifting departure time (de Palma and Fosgerau, 2011). When traffic demand (inflow) exceeds a certain threshold (the capacity of the bottleneck), a queue will start building up, and exit flow will be equal to the capacity. The queue keeps building up as long as demand exceeds capacity. At some point, demand drops below the capacity, and the queue starts to dissolve. The exit flow remains at the capacity until the queue has dissolved, after which it drops. *Figure 15* illustrates how the accumulated demand and exit flow may look, as function of time of day. Consider the average exit flow from the point in time (t') at which demand exceeds capacity to a time t'' , illustrated in *Figure 16* by the slope of the red line. During the congested period (the queue build-up and dissolution), the average exit flow is equal to capacity. After the queue has dissolved, the average exit flow drops below the capacity.

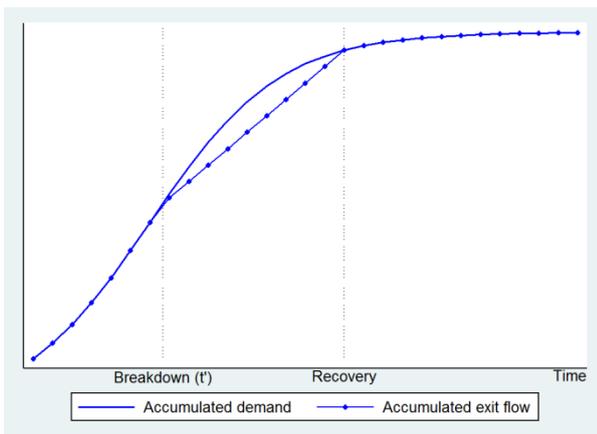


Figure 15: Bottleneck model – accumulated demand and exit flow

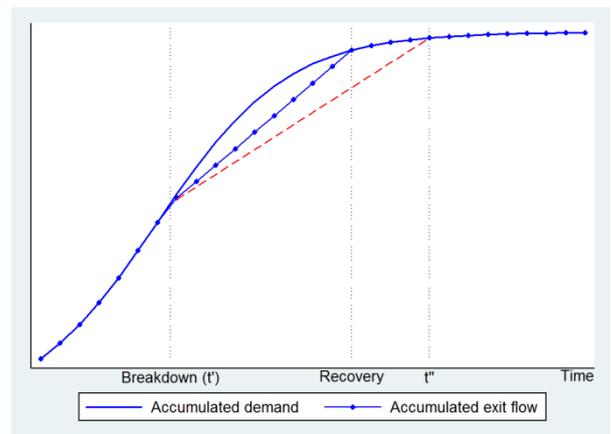


Figure 16: Inspiration to the definition of $A(t, \tau_B)$

Though the bottleneck model is a very simplified description of the congestion in our data, we expect that the average exit flow from point τ_B on has an effect on the probability of recovery, such that a lower average exit flow implies a higher probability of recovery. As stated, the flow variable F_t used in our analysis is not exactly identical to the exit flow. However we use it as an approximation and define $A(t, \tau_B)$ as the average observed flow from point τ_B on:

$$A(t, \tau_B) = \frac{ACF(t) - ACF(\tau_B)}{t - \tau_B}, \quad (3)$$

where $ACF(t)$ is the accumulated flow (F_t) at time t .

We expect that $A(t, \tau_B)$ has an effect on the probability of recovery, such that a lower value of $A(t, \tau_B)$ implies a higher probability of recovery. We can easily estimate a model including $A(t, \tau_B)$. However, information about exit flow or F_t is not available from the output of the traffic model, which predicts demand. Hence we cannot predict $A(t, \tau_B)$. So, in forecasting applications we are forced to use demand instead of F_t . What are the consequences of this? Assuming that η is the unobserved link capacity, we know from the bottleneck model that

$$A(t, \tau_B) = \min(AD(t, \tau_B), \eta), \quad (4)$$

where $AD(t, \tau_B)$ is the average demand from point τ_B on. This is because average demand exceeds the capacity η until the queue has dissolved, and afterwards it equals average exit flow. We assume that eq. (4) also holds as an approximation for the links in our analysis. Hence, in forecasting applications, we should ideally replace $A(t, \tau_B)$ by $\min(AD(t, \tau_B), \eta)$. Replacing instead by $AD(t, \tau_B)$ implies that the computed recovery probability will tend to be a little too low in time intervals with very high average demand. However, since the recovery probability is already very low in such periods, we do not consider this a major problem.

3.4 Travel Time Model

As stated above, we assume that travel time has distribution function Φ_u in the uncongested state and distribution function Φ_c in the congested state. We do not specify the shape of these distributions, as we are only interested in their means and variances. We define the random variable S to indicate the traffic state:

$$S = \begin{cases} 0, & \text{if uncongested} \\ 1, & \text{if congested} \end{cases}$$

We estimate the following four state-conditional means and variances by simple sample averages:

$$\mu_u = E(T|S = 0)$$

$$\sigma_u^2 = E((T - ET)^2|S = 0)$$

$$\mu_c = E(T|S = 1)$$

$$\sigma_c^2 = E((T - ET)^2|S = 1)$$

Given a demand profile $\{F_t\}_{t=300,\dots,720}$ we can simulate the breakdown and recovery times numerically, and for each time interval t compute the probability $p_t = P(S_t = 1)$ that traffic is in the congested state. Given the state-conditional means and variances of travel time, we can then compute the mean and variance in time interval t as:

$$\mu_t \equiv ET_t = E(E(T|S_t)) = E((1 - S_t)\mu_u + S_t\mu_c) = (1 - p_t)\mu_u + p_t\mu_c \quad (5)$$

$$\begin{aligned} \text{var } T_t &= E(E((T - ET)^2|S_t)) = E(E((T^2 + \mu_t^2 - 2T\mu_t)|S_t)) \\ &= (1 - p_t)(\sigma_u^2 + (\mu_u - \mu_t)^2) + p_t(\sigma_c^2 + (\mu_c - \mu_t)^2) \end{aligned} \quad (6)$$

We note that in this simple model where traffic has only two possible states, each with constant mean and variance of travel time, the mean travel time in any time interval t is bounded between μ_u and μ_c (since p_t is bounded between zero and one). The variance is bounded from below by the minimum of σ_u^2 and σ_c^2 , while its maximum value depends on both σ_u^2 , σ_c^2 and $(\mu_c - \mu_u)^2$.

It is possible to derive a simple relationship between the mean delay and the variance of travel time. The functional form of this relation is a direct consequence of the assumption that traffic has only two possible states, each with constant mean and variance of travel time. While we do not use this relationship in our prediction model, it is relevant to be aware of when comparing our prediction approach to international practice. As already mentioned, we assume that $\mu_u \leq \mu_c$, i.e. that the congested regime has higher mean travel time than the uncongested regime. Hence we define the mean delay in time interval t as the difference between the actual mean travel time μ_t and mean travel time in the uncongested regime μ_u . Since μ_t is bounded between μ_u and μ_c , the maximum mean delay allowed in our model is $(\mu_c - \mu_u)$. Using eq. (5), we see that the mean delay is given by:

$$d_t \equiv \mu_t - \mu_u = p_t(\mu_c - \mu_u)$$

Using this together with eq. (6) yields (for $0 \leq d_t \leq \mu_c - \mu_u$):

$$\text{var } T_t = \sigma_u^2 + d_t \left(\frac{\sigma_c^2 - \sigma_u^2}{\mu_c - \mu_u} + \mu_c - \mu_u \right) - d_t^2 \quad (7)$$

Hence, in our model, the travel time variance (per kilometer) is a concave function of mean delay (per kilometer). The standard deviation is also a concave function of the mean delay.

If we know the predicted mean delay, we can use eq. (7) to compute the travel time variance our model would predict. Note however, that applying eq. (7) on mean delays stemming from the LTM would not yield the same results as using our full model to first simulate the values of p_t over the morning and then applying eqs. (5) and (6). This is because the LTM (version 2.0) computes mean travel times (and hence mean delays) based on speed-flow curves as opposed to our approach.

4. Empirical analysis

We used the software packages Stata and Biogeme to analyse the data and estimate the parameters of the breakdown and recovery models.

4.1 Definition of link types

We distinguish between three link types, defined by the number of lanes. The types are listed in Table 5. Note that only about two-thirds of link 8 (in terms of length) has two lanes, while the last third has three lanes. However, link 8 is the closest we get to a two-lane road in our analysis sample.

Table 5: Definition of link types

Link type	Data
2 lanes	Link 8
3 lanes	Links 1,2,3,6
4-5 lanes	Links 4, 5

4.2 Definition of breakdown and recovery times

To make sure traffic starts out in the uncongested state each day in our analysis sample, we define the analysis start time to be the last 15-minute interval in the LTM night time band (9PM – 5AM), i.e. the first time interval in the analysis is 4:45AM-5:00AM, which is indexed by end time $t=300$. We define that traffic is always in the uncongested state in this first time interval, such that breakdown can happen no earlier than at $t=300$ (5:00AM).

To define the uncongested and congested states empirically, we looked at the mean and standard deviation of travel time as a function of time of day. Based on plots as *Figure 7-Figure 13*, we defined congestion to be when travel time exceeds 0.7 minutes/km (~85.7km/h) for a period of at least 30 minutes. This is a somewhat arbitrary threshold supposed to reflect that travel times in early mornings and at midday appears to be more or less constant (the variability over days is very low), and that the level of early morning and midday travel time appears to be somewhere between 0.5 and 0.7 (it varies slightly over links).

For a given day, we define breakdown to occur at the end of the 15-minute interval immediately before travel time for the first time increases to a level above 0.7 minutes/km and remains above this level for at least 30 minutes. If this does not happen, the day does not have a congested state. Given that a breakdown occurs, we define recovery to occur at the end of first 15-minute interval after breakdown where travel time is above 0.7 but drops below 0.7 in the following interval and either i) stays below 0.7 for at least 30 minutes, or ii) has already been below 0.7 once during the last hour. We demand that breakdown and recovery – if they occur – occur before noon, and so delete the few days where this is not so.

This definition implies that travel time can (temporarily) be below 0.7 during the congested period, but only for a single 15-minute interval within each hour. Note also that since a congested state must last at least at least 30 minutes, recovery cannot happen before at least

30 minutes after breakdown, implying that the probability of recovery in eq. (2) must be zero in the first 15-minute interval after breakdown.

Finally, we note that the definition of breakdown and recovery times allows for more than one congested period during the day. We are not interested in adding this complexity to the model and so decided to exclude days with multiple peaks from our analysis (29 link-days in total).

See *Figure 17* for an illustration of breakdown and recovery times.

Table 6 lists the resulting sample sizes. As link 6 has only 5 days with breakdowns, we do not include it in the analysis.

Table 6: Sample sizes

Data	Link type	Sample size (breakdown model)	Sample size (recovery model)	Days with breakdown
Link 1	3 lanes	1653 obs / 112 days	532 obs / 70 days	70
Link 2	3 lanes	1512 obs / 129 days	805 obs / 97 days	97
Link 3	3 lanes	1575 obs / 148 days	958 obs / 118 days	118
Link 4	4 lanes	4569 obs / 168 days	56 obs / 16 days	16
Link 5	5 lanes	5017 obs / 190 days	143 obs / 30 days	30
Link 6	3 lanes	2879 obs / 103 days	18 obs / 5 days	5
Link 8	2-3 lanes	2420 obs / 122 days	303 obs / 64 days	64

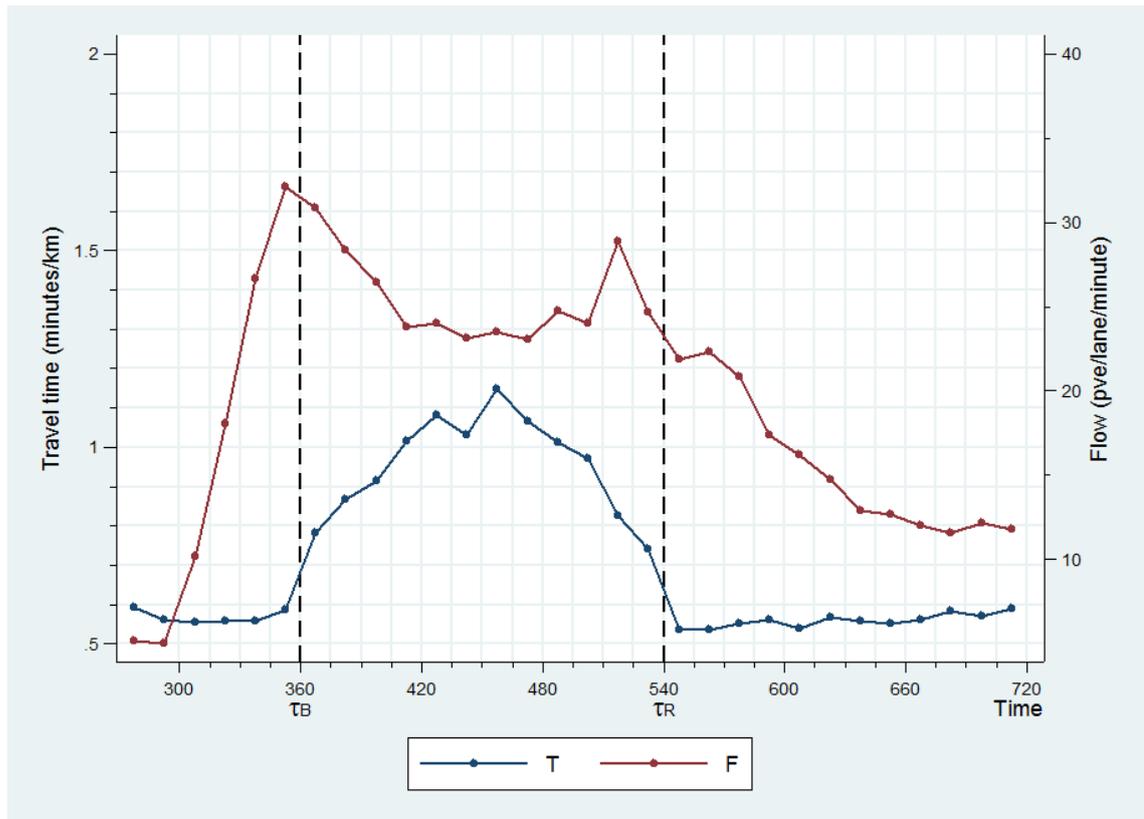


Figure 17: Illustration of breakdown and recovery times

4.3 Estimation of breakdown model

We first looked at nonparametric estimates of the breakdown probability as a function of flow F_t (Figure 18). The nonparametric estimate is a local constant regression over all observations for which breakdown has not yet occurred. The regression uses an Epanechnikov kernel (Stata default) with bandwidth 0.1. The bandwidth is deliberately chosen to be a rather small so as not to disguise variation in data by over-smoothing the curve. The nonparametric estimates reveal that the probability of breakdown increases with F_t , for values of F_t above a certain point. This is in line with expectation.

We applied the duration model in eq. (1) with f being a logistic function of F_t :

$$P_B(t) = P(\tau_B = t | \tau_B \geq t) = \frac{1}{1 + \exp(-\beta_0 - \beta_1 F_t)} \quad (8)$$

The parameter β_1 should be positive. With this model, the probability increases with F_t from zero to one following an S-shaped curve. The parameter β_1 determines how steeply the probability rises. The logistic probability model allows breakdowns to occur with very small probabilities with low traffic volumes (due to e.g. accidents), which fits well with the observed data.

We estimated the parameters β_0 and β_1 using Maximum Likelihood Estimation, with the software package Biogeme. The results are shown in Table 13 in the Appendix. Initially, we estimated different parameters for each link and link type. We are not able to estimate a breakdown model for link 5: Clearly, we observe too few days with flows high enough to cause a traffic breakdown. We note that the models for the different links are somewhat different, but since we do not have enough links to establish whether this is due to the number of lanes or other specific link characteristics, we chose to pool the data and use a single set of parameters for all link types. We stress that the number of lanes in principle still enters the model, as the flow variable F_t is measured in pce/lane/minute. Moreover, since link 3 yields counter-intuitive results in the analysis of the recovery model and links 4 and 5 have too few breakdowns to estimate a recovery model, we decided to limit the analysis to links 1,2, and 8 (cf. the discussion in section 4.4).

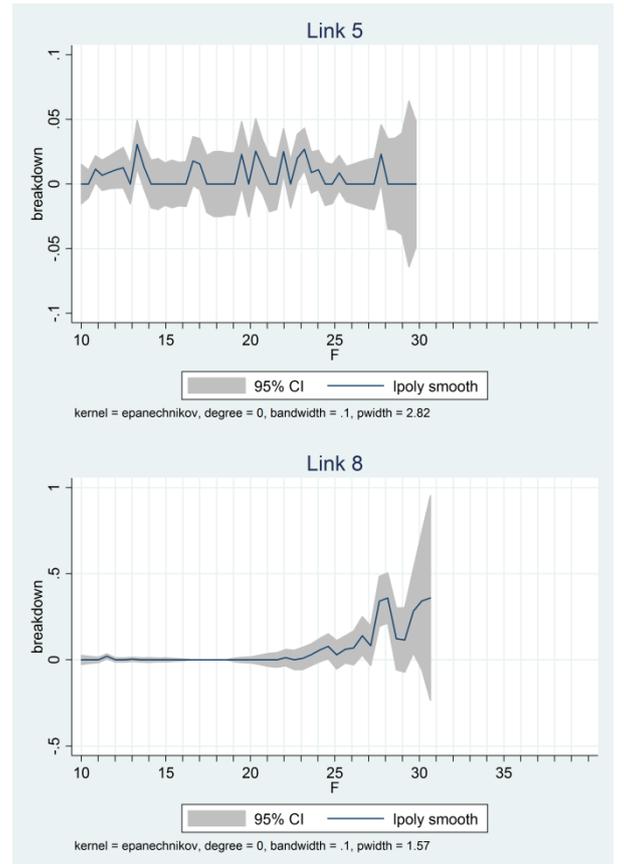
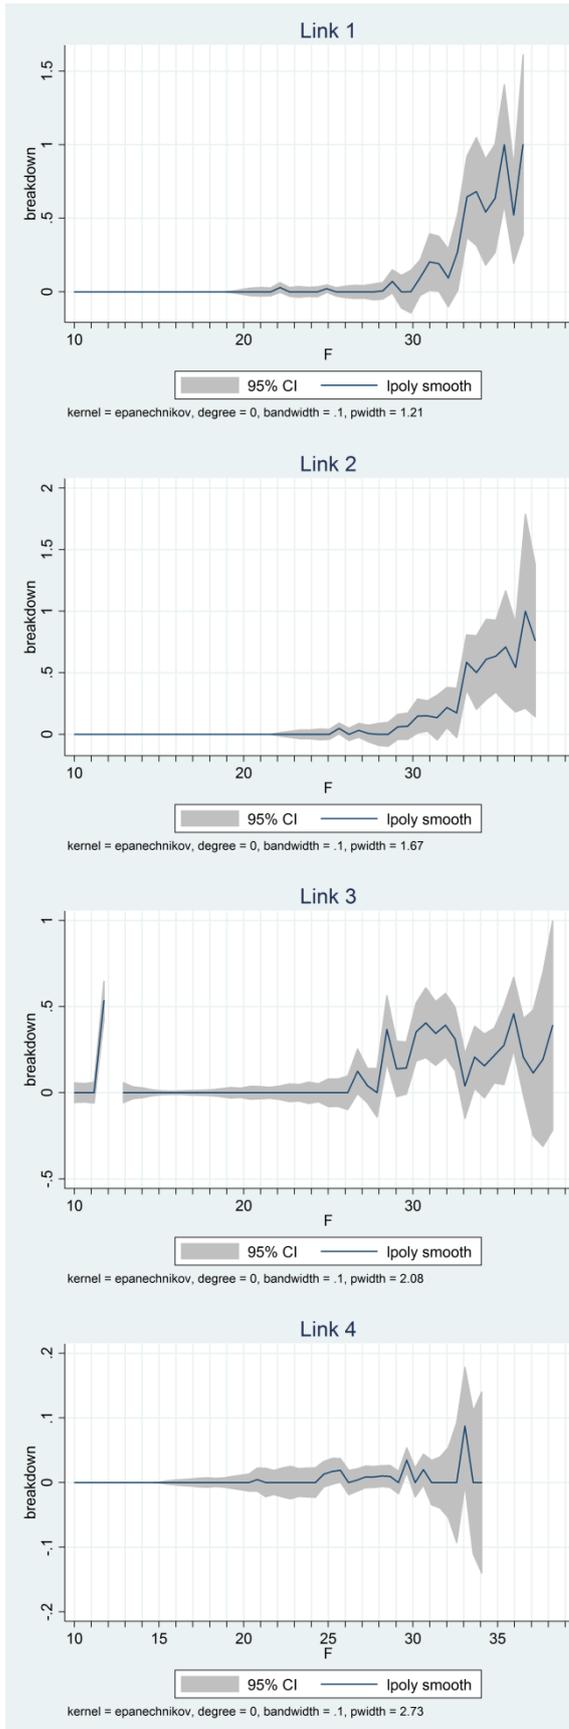


Figure 18: Nonparametric estimate of breakdown probability as function of F (by link)

Our preferred models are listed in Table 7 below. These are the specifications applied in the prediction model. The estimated $P_B(t)$ as a function of F_t is shown in Figure 19.

Table 7: Preferred breakdown model specifications

Link type	Preferred model
2-3 lanes	Model (eq. (8)) with $\beta_0 = -13.69$ and $\beta_1 = 0.3995$
4-5 lanes	None

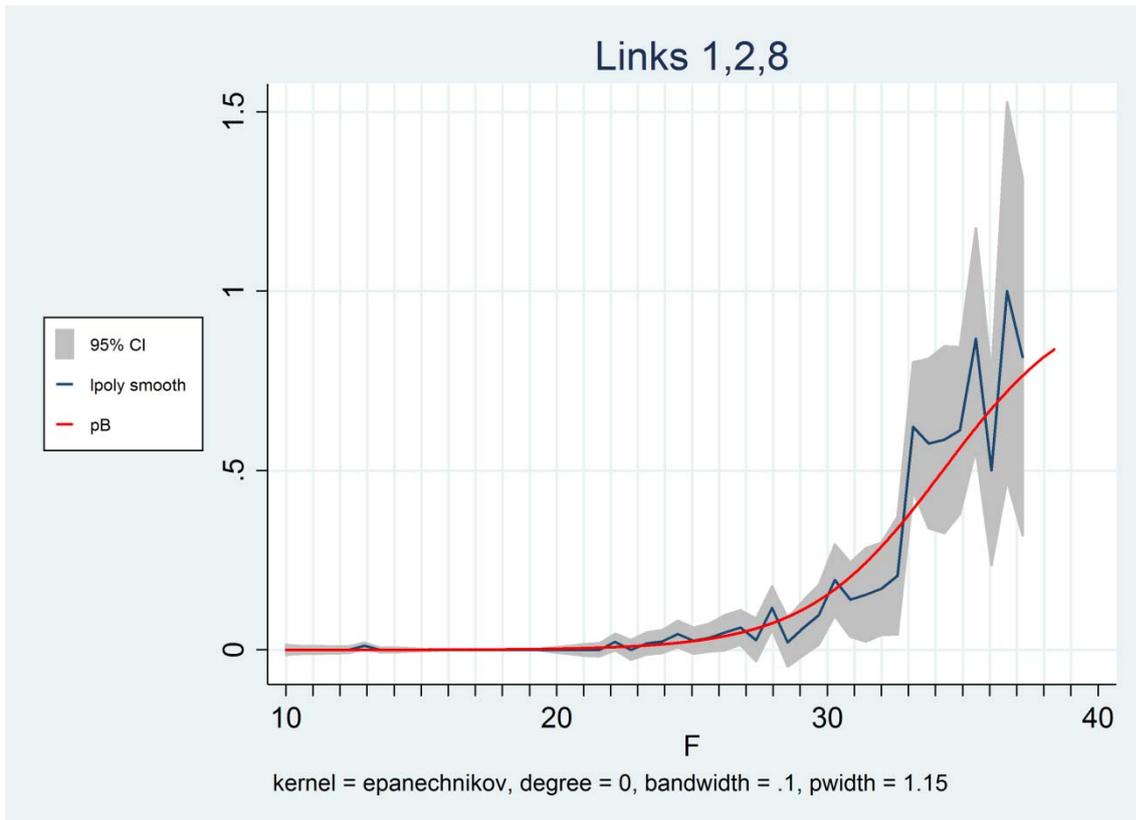


Figure 19: Estimated breakdown probability as function of F (preferred formulation, based on links 1, 2, 8)

4.4 Estimation of recovery model

To gain intuition, we first considered nonparametric estimates of the recovery probability as a function of $A(t, \tau_B)$ (Figure 20). The nonparametric estimate is a local constant regression over all observations for which breakdown has occurred but recovery has not yet occurred (except the first period after breakdown, for which the recovery probability is zero per definition of the model states, as the congested state is defined to last at least 30 minutes). Again, the regression uses an Epanechnikov kernel with bandwidth 0.1.

Figure 20 reveals an interesting feature: For links 1,2,3,8, the recovery probability is close to zero for values of $A(t, \tau_B)$ around 20-23 (approximately). For higher values of $A(t, \tau_B)$, the probability is positive and decreasing in $A(t, \tau_B)$, as theory predicts (actually, for link 3, it is not entire clear if it is decreasing, increasing or flat). A potential explanation of this pattern is that there are different types of congested states in our data: One with a clear peak in travel time, high flow and high values of $A(t, \tau_B)$, and another also with a clear peak in travel time but relatively lower flow and low values of $A(t, \tau_B)$. The latter type is often associated with very high levels of travel time, and this together with the low flow suggests that this type of congested state is related to the link capacity temporarily being reduced, due to some incident blocking parts of the lanes.¹⁰

For a given day with constant capacity, we do not believe that the recovery probability should suddenly drop to zero or a low value as $A(t, \tau_B)$ decreases. We are therefore not interested in incorporating this feature in our prediction model, as this model does not explicitly model incidents blocking parts of the lanes. So we chose to estimate separate functional forms for values of $A(t, \tau_B)$ above and below a threshold κ , and then, when predicting travel time variability, we use solely the functional form for values values of $A(t, \tau_B)$ above κ .

We estimated the recovery model (eq. (2)) using a piece-wise logistic specification with either

$$P_R(t|\tau_B = t_B) = P(\tau_R = t|\tau_R \geq t, \tau_B = t_B)$$

$$= \begin{cases} \gamma_0 & , \text{ for } A(t, \tau_B) \leq \kappa \\ 1 - \frac{1}{1 + \exp(-\gamma_1 - \gamma_2 A(t, \tau_B))} & , \text{ for } \kappa < A(t, \tau_B) \end{cases} \quad (9)$$

or

$$P_R(t|\tau_B = t_B) = P(\tau_R = t|\tau_R \geq t, \tau_B = t_B)$$

$$= \begin{cases} \gamma_0 & , \text{ for } A(t, \tau_B) \leq \kappa \\ 1 - \frac{1}{1 + \exp(-\gamma_1 - \gamma_2 \ln A(t, \tau_B))} & , \text{ for } \kappa < A(t, \tau_B) \end{cases} \quad (10)$$

As mentioned, κ is a threshold parameter. The parameter γ_0 should be positive and less than one, while the parameter γ_2 should be positive (coresponding to a negative effect of $A(t, \tau_B)$ on the recovery probability. In this model, the probability $P_R(t|\tau_B = t_B)$ is constantly equal to λ_0 for values of $A(t, \tau_B)$ below κ , and decreases following an inverse-S-shaped curve for values of $A(t, \tau_B)$ above κ .

Note that the pattern for the link 4 is very unclear and does not reveal any systematic relation between the recovery probability and $A(t, \tau_B)$ (cf. Figure 20). We therefore decided not to estimate a recovery model based on these data. We furthermore excluded link 5, since we were unable to estimate a breakdown model for this link.

¹⁰ In a few cases, we are able to confirm this hypothesis using the Trafikman data. For the remaining days there is no reporting of incidents. However, according to the Road Directorate, there may be unreported (small) incidents that could have similar effect, such as a car stopped in the emergency lane or dropped items on the road.

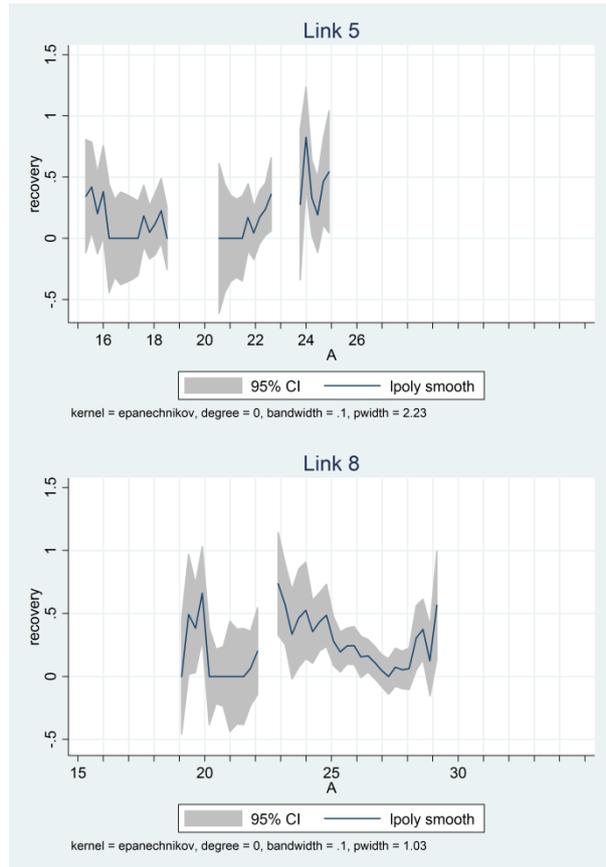
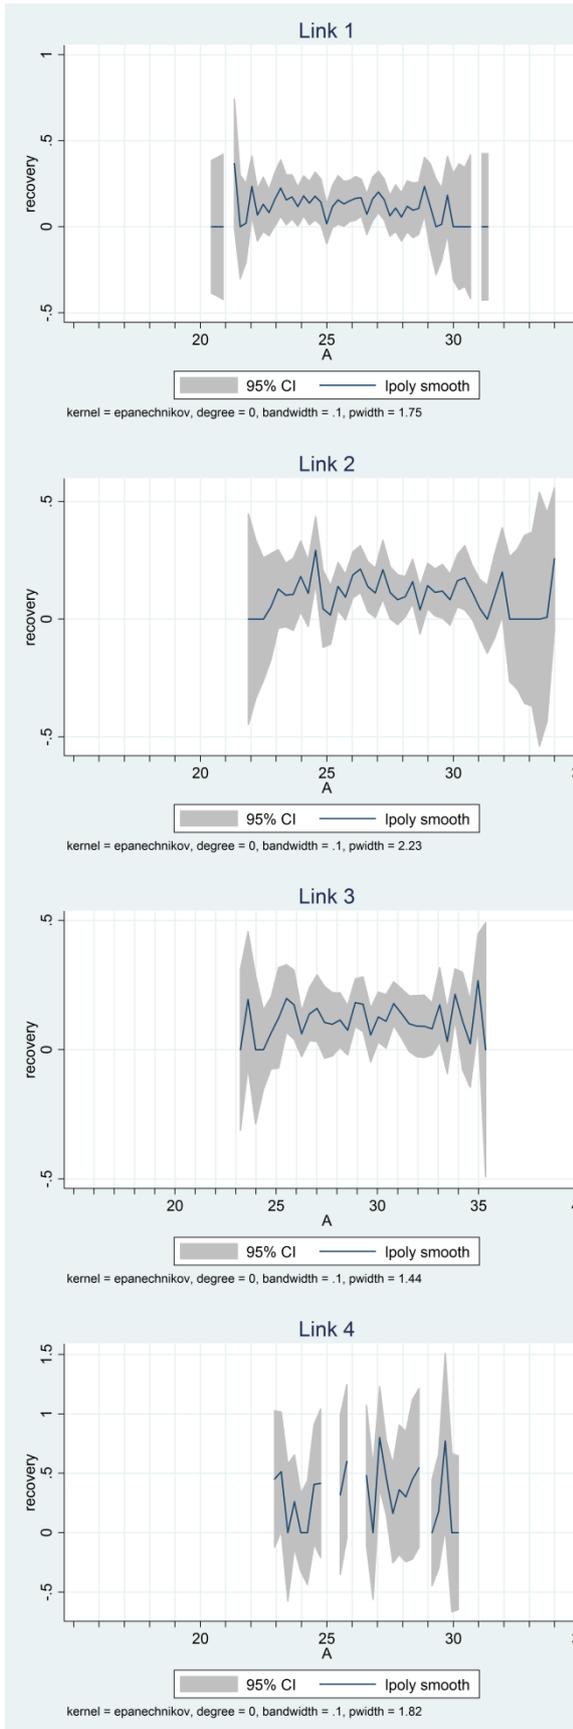


Figure 20: Nonparametric estimate of recovery probability as function of A (by link)

We estimated the parameters γ_0 , γ_1 and γ_2 using Maximum Likelihood Estimation, with the software package Biogeme. As mentioned above, we prefer a common model specification across link types, but as a data check we estimated different parameters both for each link separately and for each link type as well as for the combined sample. We estimated both the model in eq. (9) and the model in eq. (10) for four different values of κ : 20, 21, 22 and 23. All results are shown in Table 14 in the Appendix.

First, we notice that the link-specific models for links 1, 2, 3 are not always identified. The problem lies in identifying γ_0 : In the link-specific models, there are too few observations with $A(t, \tau_B) \leq \kappa$. As a check, we therefore estimated models with γ_0 fixed to zero (cf. Table 15 in the Appendix). From these results, we find that the models for links 1 and 2 behave as expected (γ_2 is consistently positive, though not significantly), while for link 3, γ_2 is consistently of the wrong sign (though not significantly). This effect from link 3 carries through when we consider a common model across links: For the sample consisting of links 1, 2, 8, the parameter γ_2 is always significantly positive, while this is rarely the case for the sample of links 1,2,3,8. This led us to prefer the common model for links 1, 2, 8.

Second, we used the maximum likelihood value to determine the best κ -value and to compare the models in eq. (9) and eq. (10). The best κ value is 23, and for this value the logarithmic formulation in eq. (10) is slightly better than the formulation in eq. (9) and does not “suffer” from insignificant parameters. We summarise the preferred results in Table 8.

The estimated recovery probabilities are shown in *Figure 21* together with the nonparametric estimates.

Table 8: Preferred recovery model specifications

Link type	Preferred model
2-3 lanes	$P_R(t \tau_B = t_B) = P(\tau_R = t \tau_R \geq t, \tau_B = t_B)$ $= \begin{cases} 1 - \frac{1}{1 + \exp(8.907 - 3.261 \ln(23))} & , \text{ for } A(t, \tau_B) \leq 23 \\ 1 - \frac{1}{1 + \exp(8.907 - 3.261 \ln A(t, \tau_B))} & , \text{ for } 23 < A(t, \tau_B) \end{cases}$
4-5 lanes	none

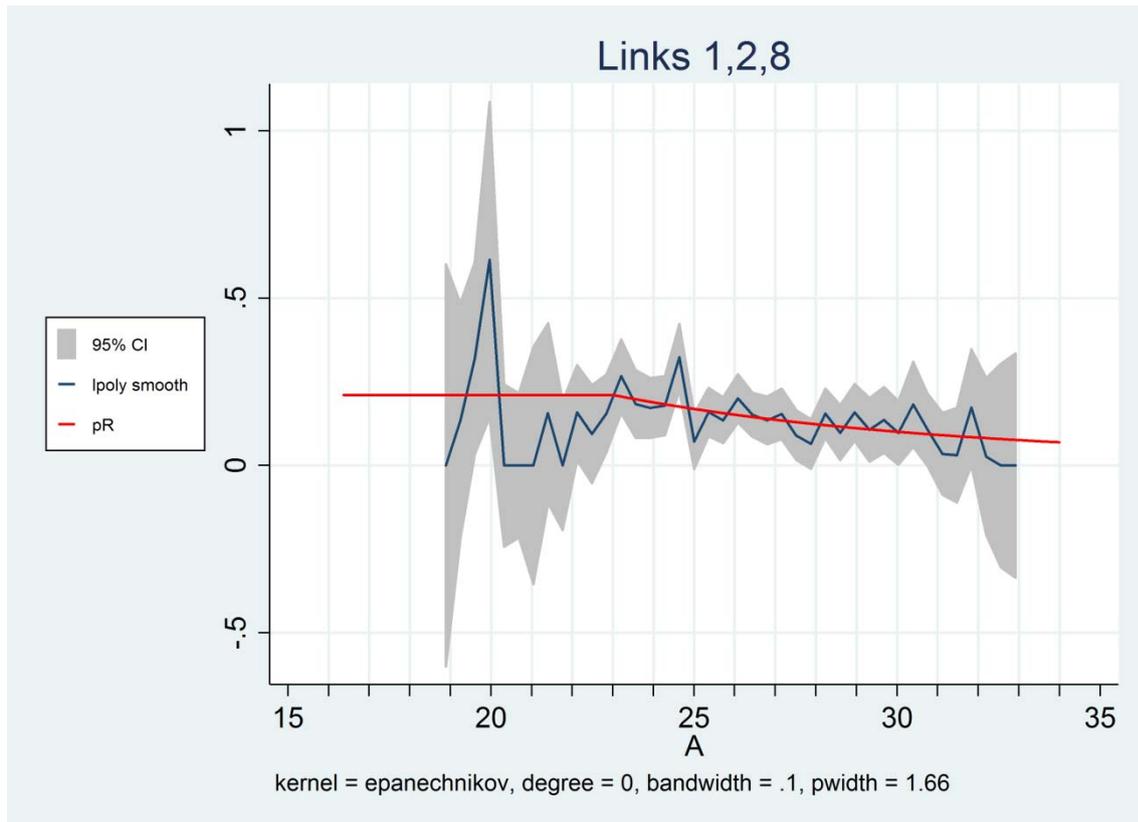


Figure 21: Estimated recovery probability as function of A (preferred formulation, eq. (10) with $\kappa=23$, based on links 1, 2, 8).

4.5 Model validation

To validate the model, we used it to simulate the breakdown and recovery times for the estimation sample (links 1, 2, 8), and compared to the pattern of actual breakdown and recovery times. For the simulation, we used the actual flow pattern over the morning period for all days in the estimation sample and generated 1000 “copies” of each day. For each day, we then used the models in Table 7 and Table 8 to predict breakdown and recovery times. Figure 22 and Figure 23 show the pattern of real and simulated breakdown times, respectively. The rightmost column in each histogram represents the days without a breakdown. The breakdown model does not reproduce this share exactly: In the simulation, the share is 42% as opposed to 36% in the real data. Apart from this, however, the breakdown model reproduces the pattern of breakdown times almost surprisingly well, considering the very simple model formulation applied.

Rather than comparing recovery times, we compare the pattern of peak durations. This is shown in Figure 24 and Figure 25 for the real and simulated data, respectively. Here, the days without a peak (i.e. without a congested period) are represented by the leftmost column with zero peak duration. The average peak duration is captured quite well: 118 minutes for the simulated data and 121 minutes for the real data. However, the spread of the simulated peak duration is much larger than of the real data (cf. Figure 25), suggesting that the predictive power of the recovery model is not so strong.

We attempted to remedy this in several ways, though our options were limited by the fact that any explanatory variable in the model would have to be available from the output of the national traffic model. We tried to capture the “multi-peaked”- shape in *Figure 24* by including information about the share of long vehicles, information about the traffic flow in the time period immediately before breakdown, and information about the value of $A(t, \tau_B)$ in the first time period after breakdown. We also tried to omit Fridays from the estimation sample. Though some of these affect the pattern of simulated peak durations, we did not find that they produced a better fit. As an example, including the share of long vehicles excludes the existence of very long peak durations, but at the cost of significantly over-predicting the number of peaks lasting 90 minutes or less.

In conclusion, we chose to stick to the simple model formulation in Table 8. Despite its lack of (predictive) power, it still serves as a simple approximation of reality.

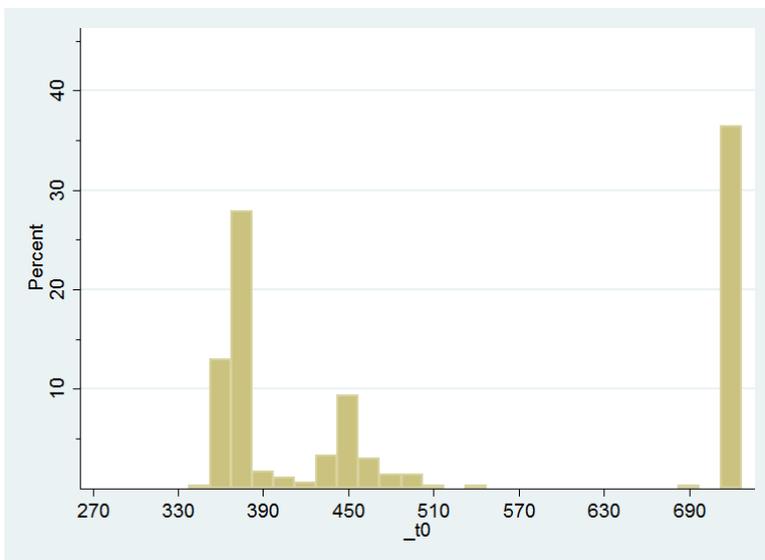


Figure 22: Distribution of breakdown times in the sample (links 1, 2, 8).

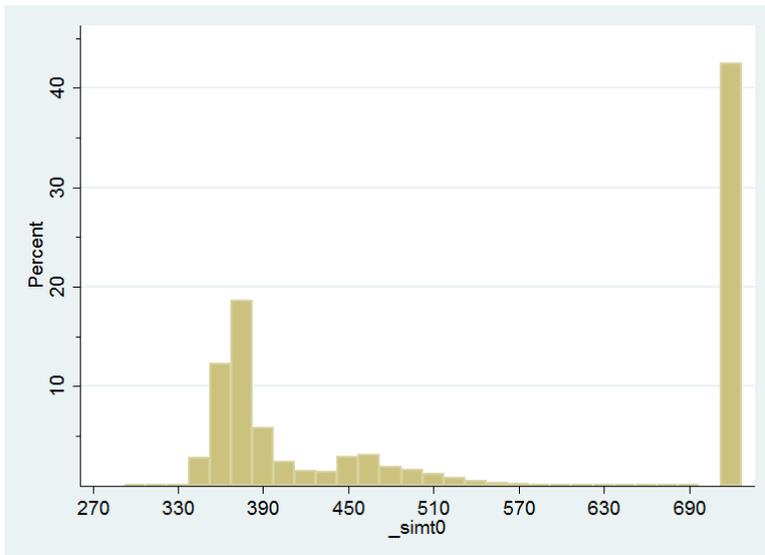


Figure 23: Distribution of simulated breakdown times.

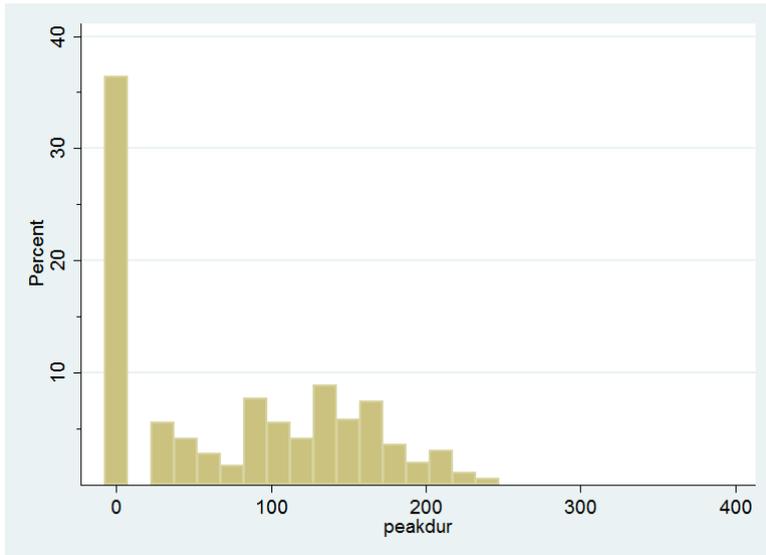


Figure 24: Distribution of peak durations (in minutes) in the sample (links 1, 2, 8).

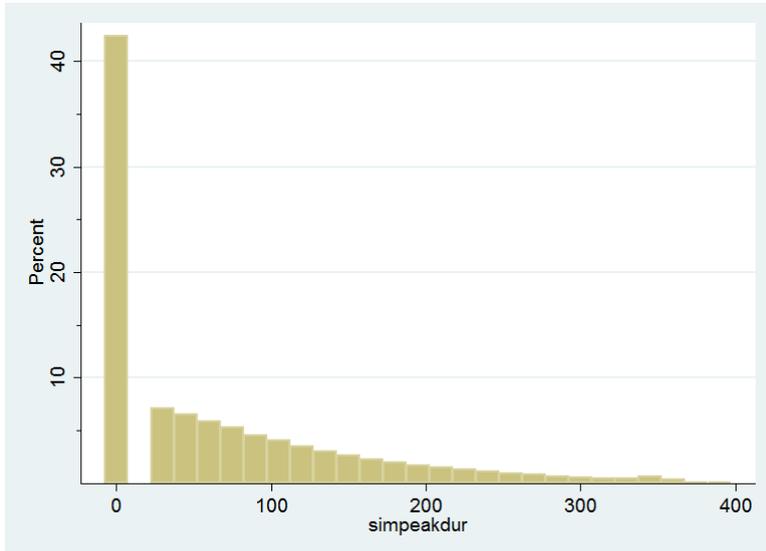


Figure 25: Distribution of simulated peak durations (in minutes).

4.6 Estimated state-dependent means and variances of travel time

For each traffic state (congested/uncongested), we computed means and variances of travel time as simple sample averages over the time periods from “Typical weekdays” in the estimation sample (i.e. links 1,2,8). Note that when we estimated the breakdown and recovery models, we used not only “Typical weekdays” but also “Special days” (cf. sec. 2.4) in order to get sufficient variation in the data. However, when forecasting travel times we used only “Typical weekdays” for consistency with the national traffic model. Table 9 shows the values, which we apply in the prediction model.

Table 9: Applied means and variances of travel time by traffic state

μ_u	σ_u^2	μ_c	σ_c^2
0.58	0.00096	1.23	0.19

5. Application

5.1 Outline of implementation

Here we describe the procedure generating travel time variability based on traffic flow output from LTM. The procedure simulates the transitions between traffic states by simulating the breakdown and recovery times using the estimated models in sections 4.3 and 4.4. It computes the probabilities p_t that traffic is in the congested state and calculates the mean travel time and its variance using eq. (5) and (6).

The algorithm below computes mean and standard deviation of travel times for a given link in a given direction, for the morning (AM) period. A similar algorithm should therefore be run for the PM-period. In the algorithm, the first 15-minute interval in the simulation is 4:45-5:00AM. The choice of this starting period is not crucial – in principle the only criteria is that it is sufficiently early to ensure that breakdowns do not occur before this time. For the application with LTM output, we therefore suggest that the starting period should be one of the last 15-minute time intervals in the LTM night time band (9PM – 5AM).

Algorithm to compute mean and standard deviation of travel time for AM-period

1. The traffic flow for the link during the day is read from an LTM run. In addition, the relevant attributes of the links (number of lanes etc.) is read and flow per lane is computed. Each link has 10 traffic flows corresponding to the 10 LTM time bands – see Table 10 below. The flow in a LTM time band is spread out to the 15 minute time intervals of the simulation model assuming a constant flow in all periods corresponding to a LTM time band.
2. The peaks are simulated (including travel time and travel time variability) using the estimated model and traffic flow ($F_{300}, F_{315} \dots F_{720}$) from step 1.
 - a) The traffic flow is made stochastic by drawing from a multinomial distribution with 10 outcomes giving the likelihood of a flow in the neighbourhood of the observed, ranging from 19% below and 18 % above. The same change is used for all periods during a day. The distribution is based on the observed variation in the flow in the data.
 - b) For each period, we compute the probability $P_B(t)$ as a function of $F_{300}, F_{315} \dots F_{720}$ using the prediction model in Table 7.
 - c) For each period t we draw from a binomial distribution with probability $P_B(t)$ of success. If a transition occurs (success), the variable B_t is assigned the value 1, otherwise $B_t = 0$.
 - d) We initialise the state variable $S_t=0$ for all t . The first period t' with $B_{t'}=1$ during the day (if any) is identified and is chosen as the time of transition from uncongested to congestion. We set the state variable $S_t=1$ for all periods $> t'$.
 - e) For each period we compute the probability of recovery $P_R(t|\tau_B = t')$ as a function of $F_{300}, F_{315} \dots F_{720}$ using the prediction model in Table 8.
 - f) For all periods t with $S_t=1$ **except the first** ($t = t' + 1$), we draw from a binomial distribution with probability $P_R(t|\tau_B = t_B)$ of success. The first time t'' for which a transition occurs (success) is identified. The state variable is changed to $S_t=0$ all periods $t > t''$. Thus $S_t=1$ indicates the time period with congestion. A day can only

have 0 or 1 period with congestion. P_R is set to 1 at noon ($t = 720$) to ensure that peak hours end at this hour at the latest.

- g) The calculations a) to f) is repeated a number of times (e.g. 1000)
 - h) For each time period t the share p_t of instances with $S_t=1$ over all the repetitions is computed. The number of repetitions is chosen such that these shares converge.
 - i) The expected mean travel time and variance is computed for each period using the shares p_t from h) using eq. (5) and eq. (6) and the estimated state-dependent travel time means and variances from Table 10.
 - j) Costs per minute per road-kilometre are calculated separately for travel time and the standard deviation taking into account the flow, the number of lanes, and the share of long vehicles. The cost per car per hour is taken from the unit prices from The Ministry of Transport.
3. Finally, the resulting travel times, variances and costs are aggregated from 15-minute time intervals to LTM's time bands. The resulting travel time means and variances are computed using simple averages (though they are probably not independent).

Table 10: LTM time bands

LTM time band	Period
1	9PM -5AM
2	5AM -6 AM
3	6AM -7AM
4	7AM -8AM
5	8AM -9AM
6	9AM -3PM
7	3PM -4PM
8	4PM -5PM
9	5PM -6PM
10	6PM -9PM

5.2 Application example

For use in our application examples, we implemented the above procedure in SAS and performed the following test computations on the subsample of typical weekdays:

5.2.1 Test of the flow.

It is checked that the way flow is made stochastic corresponds to the observed flow. This is done by replacing the static flow input with the observed average flow for each time period. It turns out that the average flow within each time period is maintained when 1000 repetitions is applied. The average relative standard deviation (of flow) during the morning is maintained as well, but there are significant differences when each time period is considered separately. The graph below (*Figure 26*) with data (incl. flow variability) for link 1+2 as an example shows how the modelled standard deviation of flow by construction follows the mean flow, while the observed standard deviation has a maximum later than the mean flow peak.

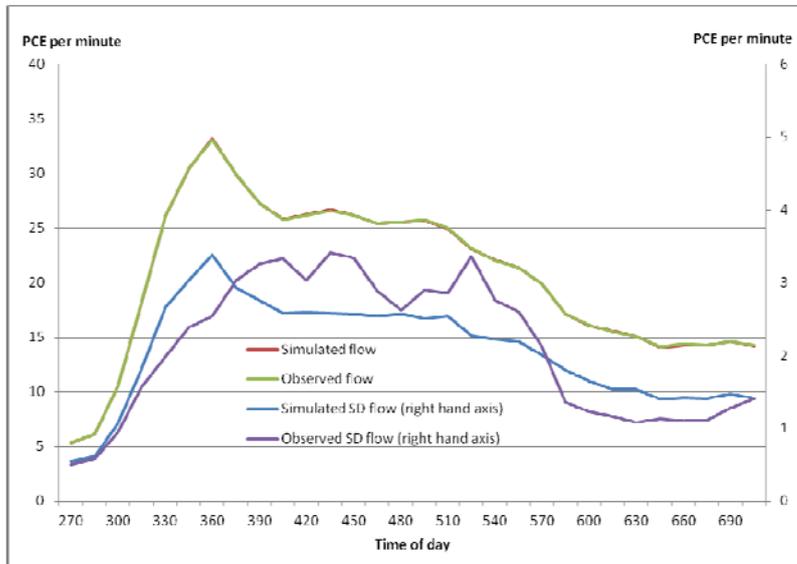


Figure 26: Simulated and observed mean and standard deviation of flow over the morning, links 1+2, 1000 repetitions

This pattern emerges for the other links as well: The observed standard deviation is approximately proportional to the mean flow, but with a certain time delay. The modelled standard deviation is very close to proportional to the mean flow and is thus higher than the observed at the start of the peak, but later in the morning the opposite is the case.

5.2.2 Test of the peak frequency and duration

The frequency of days with peak for the links 1 and 2 is 74.9% in the data. When simulating with 1000 repetitions on the data from these three links, the resulting peak may vary a little from simulation to simulation, but is typically 78-80% which is satisfactory.

In the data there is a clear pattern that there is a high share of days with peaks shorter than one hour and a high share of peaks with duration around 2½ hour. This pattern is not reproduced in the simulations since the model by construction gives lower probability to long peaks than to short peaks. This is illustrated in the histograms in *Figure 27*. The patterns here are consistent with those produced in section 4.6, that were computed based on the actual rather than simulated flow.

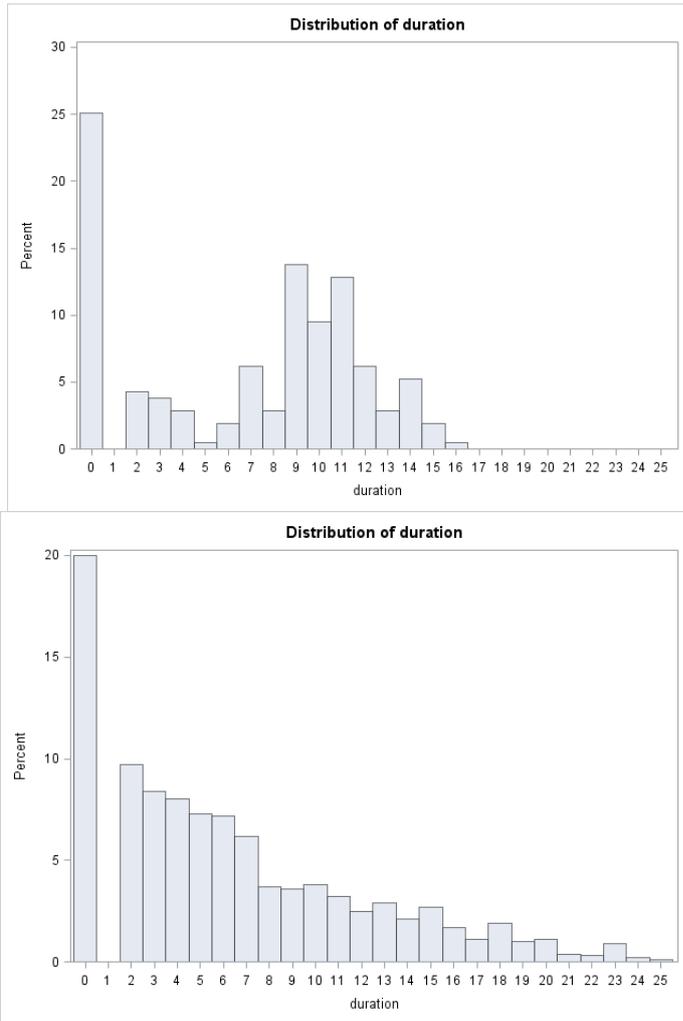


Figure 27 Peak duration (in quarters of an hour), observed (above) and simulated (below), links 1+2, 1000 repetitions

5.2.3 Test of travel time and travel time variability

The mean travel time and travel time variability for the links 1 and 2 are 0.700 and 0.313 (std. dev.) over the period 4:30AM to noon. The corresponding simulated values are 0.715 and 0.291. The mean travel time is well described, but the model seems to underestimate the travel time variability to some extent. This is partly caused by the bias in peak duration seen in Figure 27.

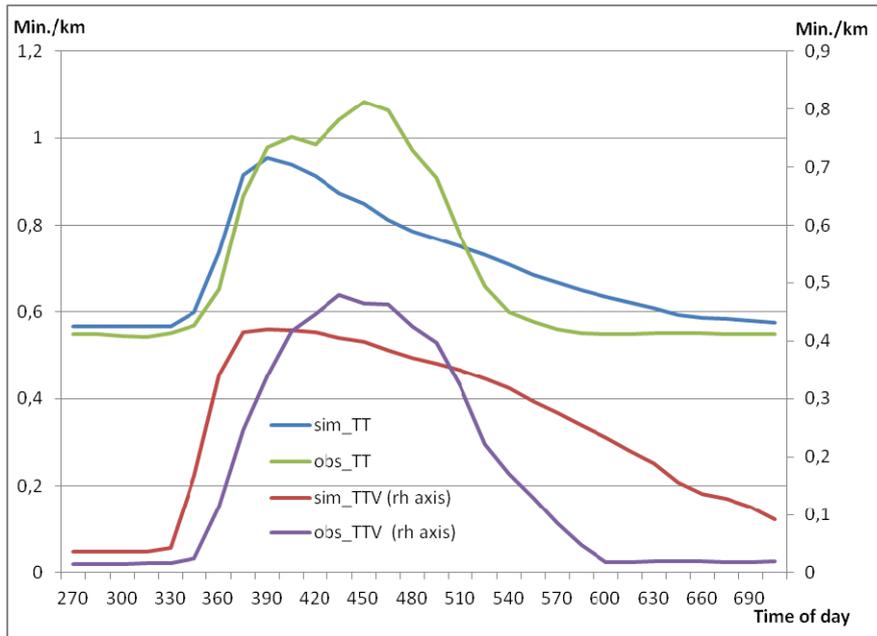


Figure 28: Observed and simulated mean travel time (TT) and TTV, links 1+2, 1000 repetitions

In Figure 28 the relation between the observed and simulated mean travel time and travel time variability is shown for each 15-minute time interval in the morning. In general, the model seems to overestimate mean travel time and TTV at the beginning and at the end of the peak hours, but underestimates them during the most congested hours. The reason is probably partly the flow simulation that overestimates the flow variation at the start of the peak and underestimates it later in the morning.

5.3 Applying the model for experiments

5.3.1 Simulated volume-delay relations

A first application of the model is to simulate mean travel times and TTV for different flow levels. We re-run the simulations from sec 5.2, scaling the input flow profile such that it is 30%, 40%, 50% 170% of the sample average flow. The resulting mean travel time and TTV (as weighted averages over the entire morning, weighted by the flow in each time period) are shown in Figure 29 and Figure 30.

As expected, both mean travel time and TTV increase with the traffic volume. For very low traffic volumes, they become constant. This is because the model then predicts an (almost) zero probability for a breakdown and thus assigns the uncongested mean and standard deviation from Table 9 to all time intervals. For very high traffic volumes, the opposite is the case: The model then predicts that breakdown occurs (almost) with certainty, and assigns the congested mean and standard deviation from Table 9 to all time intervals. The model is therefore not quite realistic for very large traffic volumes: It operates with only two states (congested and uncongested) and each is assumed to have a constant distribution of travel times independent of the traffic volume. We expect the model to underestimate both mean travel time and TTV for very large traffic volumes.

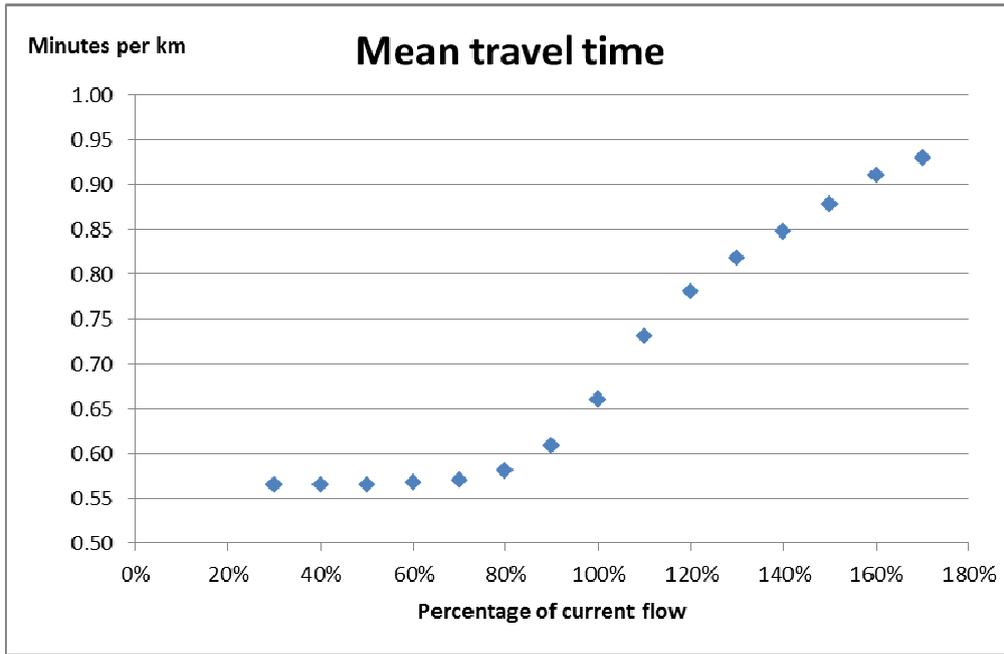


Figure 29: Simulated mean travel time (average over the period 4:30AM - noon) as function of traffic volume.

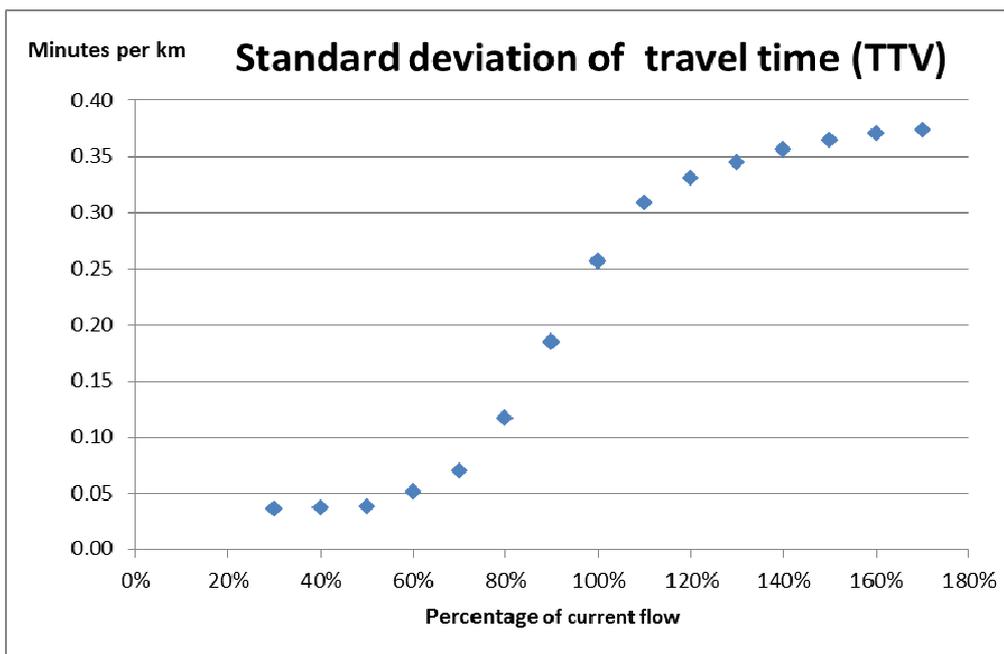


Figure 30: Simulated standard deviation of travel time (average over the period 4:30AM - noon) as function of traffic volume.

5.3.2 Applications to congestion-reduction scenarios

The model has been applied to four examples of initiatives directed at reducing congestion and travel costs. The experiments are in general carried out using the flow data for links 1 and 2 to ensure that all links have the same number of lanes. The only exception is experiment number four which is based on flow data for link 8.

The way the time costs are calculated demands some more detailed description: We use the information on vehicles length in meters (0-5.8, 5.8-12.5 and 12.5-) as indicator for the three types of vehicles:

- Passenger cars (average load factor and share of business travel)
- Vans
- Trucks

The link between length and vehicle type is simply that passenger cars are assumed to be short, vans 5.8 to 12.5 m and trucks more than 12.5 m. This is not very accurate, but a manageable simplification. The value of time for goods in vans and trucks are ignored since we have no information on this. We value both mean travel time and travel time variability (standard deviation) using a reliability ratio of 1 (i.e. one minute of travel time standard deviation equals one minute of travel time). This is based on DTU Transport, 2008.

The costs are presented as DKK-2015 per time unit per road-kilometre and cover all lanes of the road segment. The expected traffic flow and observed share of long vehicles for each time of day are used in the calculations.

The applied time values are taken from the Danish Unit Prices for cost-benefit analysis and are:

- Passenger cars: 186 DDK per vehicle per hour
- Vans: 375 DDK per vehicle per hour
- Trucks: 518 DDK per vehicle per hour

We consider four congestion reducing initiatives. The scenarios we consider are meant as illustrations of the model, and should not be interpreted as more than this. The assumptions we make regarding implementation of the initiatives are highly simplified and not completely realistic. For example, we assume throughout that the total traffic demand is unchanged (and equal to that observed in our data), while a more realistic calculation would demand that new traffic demand profiles were computed using a traffic model. The four initiatives are defined as follows:

1. **Peak spreading.** This could be the result of road user charging differentiated according to time. The flow is smoothed to be maximum 25 pce's per lane per minute. The excess flow is moved as little as possible to time bands before and after the flow peak in a 50-50 split. This is to mimic the effect of a time varying road user charge with high charges during the peak and low otherwise.
2. **Ramp metering.** The probability of breakdown (p_B) is reduced by 20% in all time periods to reflect the effect of ramp metering.

3. **Expanding the road from 3 to 4 lanes, keeping demand fixed** – flow per lane is reduced to $\frac{3}{4}$ of its current level.
4. **Expanding the road on link 8 from 2 to 3 lanes, keeping demand fixed** – flow per lane is reduced to $\frac{2}{3}$ of its current value.

The effect from the smoothing (initiative 1) can be seen from *Figure 31*.

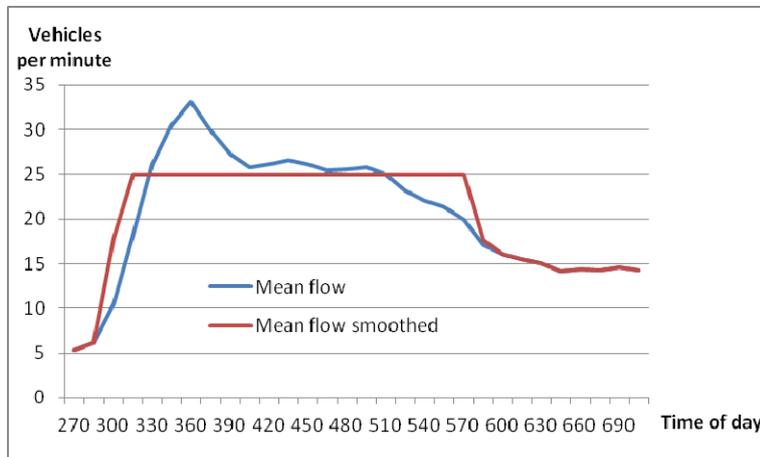


Figure 31: Average flow before and after smoothing, links 1 and 2.

To ensure convergence 10,000 repetitions have been applied in the simulations and the results from the four initiatives are summarized in Table 11, which shows the aggregated travel time costs per morning per road-km. For comparison with current practice, Table 12 presents the travel time costs that would be attributed to free flow travel time and mean delay, and their sum.

Table 11: Main results from the model simulations (4:30AM - noon)

Experiments	Share of peak days	Mean peak duration	Cost of mean TT	Cost of TTV	Total travel time cost
Initiative		No. of 15-minute intervals	DKK/morning/road-km	DKK/morning/road-km	DKK/morning/road-km
Base case (links 1 & 2)	88.9%	8.08	68,718	26,784	95,501
1. Peak spreading ($F \leq 25$)	67.8%	6.45	61,350	23,633	84,982
2. p_B reduced by 20%	82.6%	7.90	67,081	26,150	93,231
3. From 3 to 4 lanes	14.5%	5.83	53,928	11,358	65,287
Base case (link 8)	52.4%	6.41	32,230	10,696	42,926
4. From 2 to 3 lanes	1.7%	6.20	28,021	2,813	30,834

Comparing the four experiments (Table 11) reveals that user charging and ramp metering seem to have moderate effects on time related travel costs, whereas investments in additional lanes seem to be stronger tools. Clearly, this should not be taken to imply that the latter initiatives are better, as we focus on benefits in travel time costs only and consider neither further benefits nor the size of the investment needed. Again, we stress that the shown experiments should be considered only as examples of the application possibilities that the model has: The concrete implementation of the initiatives is much too simple to allow realistic analyses and conclusions.

Table 12: Travel costs based on current practice (4:30AM - noon)

Experiments	Cost of free flow time	Cost of mean delay	Total travel time cost
Initiative	DKK/ morning/ road-km	DKK/ morning/ road-km	DKK/ morning/ road-km
Base case (links 1 & 2)	51,958	24,797	76,756
1. Peak spreading ($F \leq 25$)	51,958	14,063	66,021
2. p_B reduced by 20%	51,958	22,376	74,334
3. From 3 to 4 lanes	51,958	2,915	54,874
Base case (link 8)	27,890	6,423	34,313
4. From 2 to 3 lanes	27,890	195	28,085

Comparing Table 11 and Table 12 shows that in these simple scenarios, the current approach of computing travel costs based on free flow travel time and mean delay would greatly underestimate the costs, but would result in a similar ranking of the initiatives. For initiatives 1-3, the relative gain in total cost compared to the base case differs only slightly between Table 11 and Table 12, while for initiative 4, the difference is somewhat larger.

Below in *Figure 32*, *Figure 33* and *Figure 34*, we present more detailed information about the four initiatives, in terms of their effect on mean travel time, travel time variability and total travel time costs in each of the 5 relevant LTM time bands, i.e. 5-6, 6-7, 7-8, 8-9 and 9-12.

From the figures it can be seen that including travel time variability may change the estimated effects of various initiatives to mitigate congestion. Often the variability adds around 50% to the estimated costs and in the cases where the road is expanded the change in variability may account for more than half of the economic gains related to travel time.

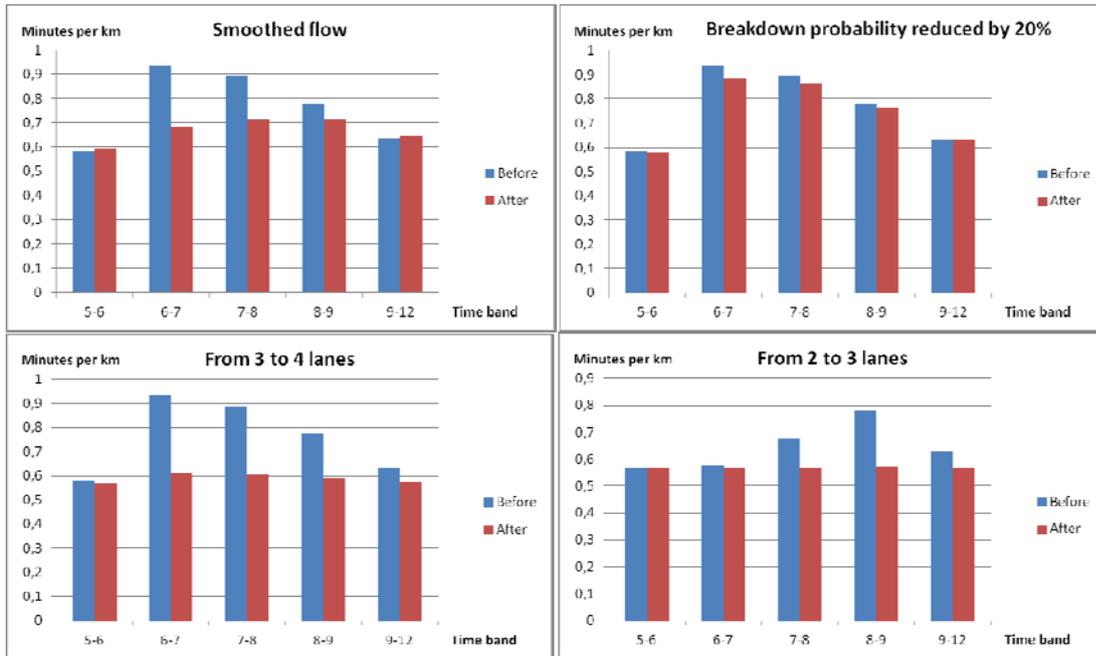


Figure 32: Simulated mean travel time before and after initiative

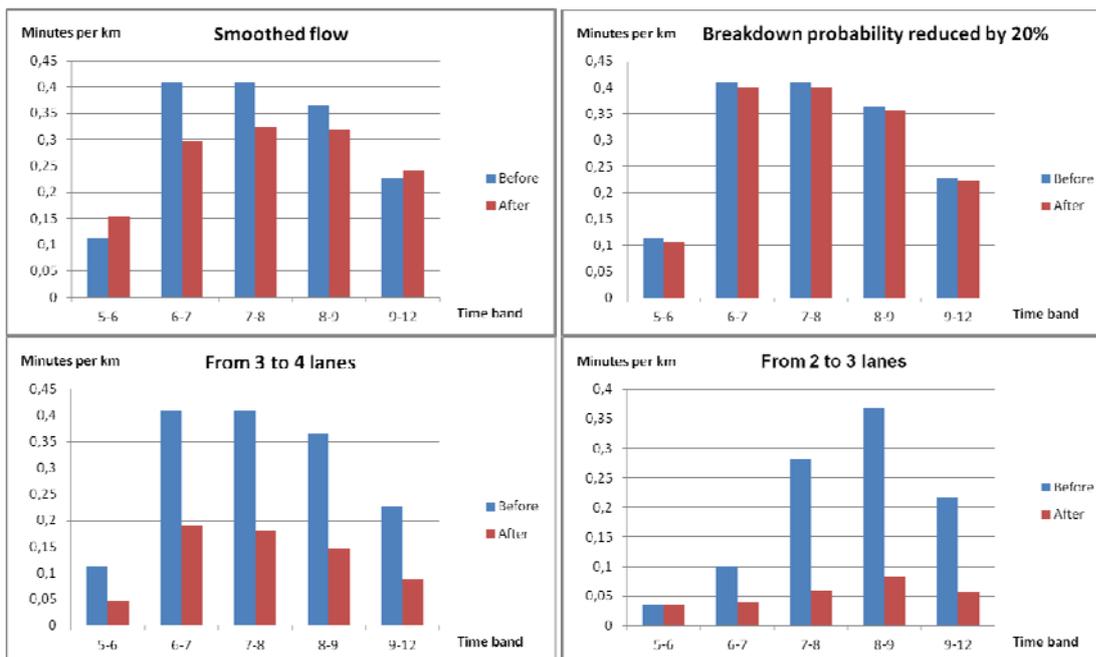


Figure 33: Simulated standard deviation of travel time before and after initiative

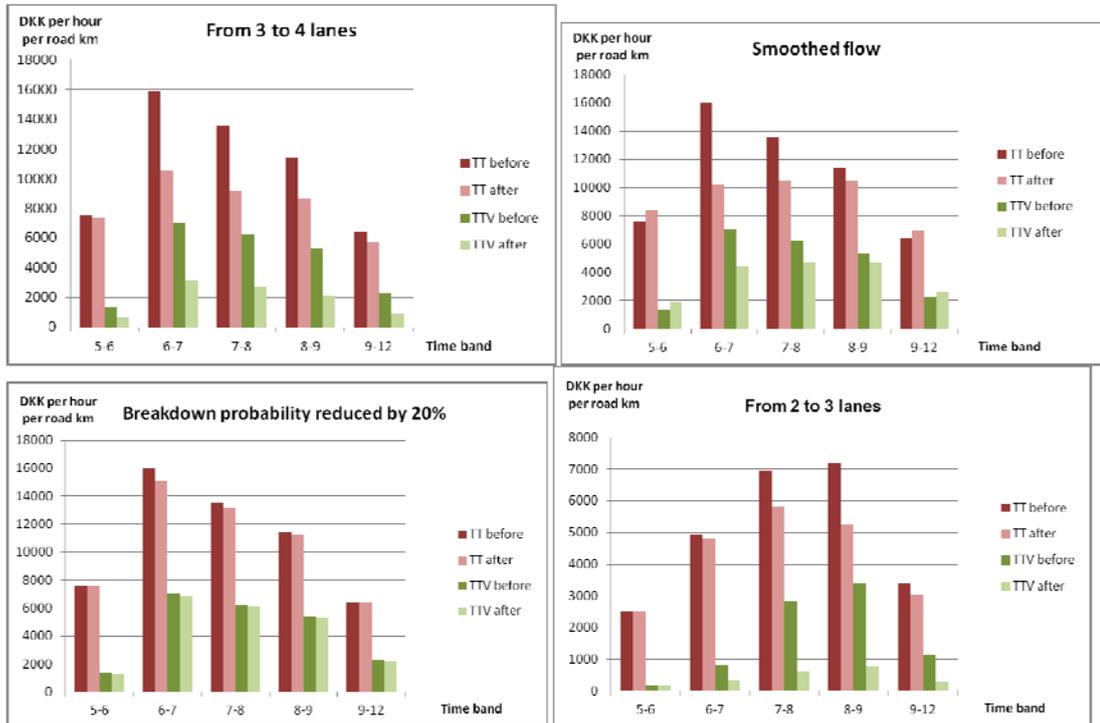


Figure 34: Simulated costs from mean travel time and travel time variability before and after initiative

6. Conclusions and recommendations for future model development, implementation and use in traffic model

In this section we discuss the future model development and data collection that we find necessary to achieve the goal of a simple forecasting approach that is ready to implement and use with the LTM.

6.1 Handling spillback effects

While we are confident that we have succeeded in developing a model that takes account of the dynamic nature of congestion and the potential endogeneity issues related to modelling travel times and traffic flows, we have at this point not succeeded in controlling for spillback effects. Given the nature of our data, that represent a series of short adjacent motorway links from a road with multiple potential bottlenecks (the entry merging ramps), this remains an important issue that should somehow be handled before the method can be implemented.

We believe it is possible to get closer to this goal with the current dataset set up for this project, even though we have not succeeded doing this within the project due to time limitations. It may not be possible to obtain a theoretically satisfactory modelling of bottlenecks and spillbacks due to data limitations (no data from ramps available) and the aggregated nature of the data (link level travel times). However, it is likely possible to approximate spillback effects to some degree and check if this improves the predictive performance of the recovery model. A more detailed analysis may demand better data in terms of observations of traffic on relevant ramps.

It may be worth seeking inspiration in an ongoing research project about modelling of spillbacks in static route choice models, led by Christian Overgård Hansen at DTU Transport (Overgård Hansen et al., 2014). The objective of the project is a practically applicable approach that can be applied with the existing static route choice models to make up for the fact that these static models cannot handle spillback effects. Their approach uses the simple deterministic queue model described in The Danish Road Directorate (2010a) to compute mean delays caused by queuing at bottlenecks. This model computes the area between the accumulated demand curve and the accumulated exit flow curve, and thus has some similarity with our recovery model. At this point, their approach does not suggest how to model TTV (which is not the point of their analysis), and compared to our approach it has the drawback that it does not handle dynamic effects.

6.2 Generalising to other motorways and remaining road network

In section 2 we discussed our (rather strict) demands on the data we applied in the analysis. Still, in section 4 it turned out that even these data were not sufficient to allow estimation of the model parameters for all links in the sample.

Due to the extensive data demands to estimate the model, we believe that it is not realistic in the short run to develop separate models for all road types. We suggest prioritizing models for different types of motorway sections (more than in our analysis) and other larger roads where the data are already available. For some motorways, i.e. the Helsingør Motorway, the existing data would be sufficient. The same might apply to few of the other greater roads, while we believe that the communal roads are not sufficient covered to estimate a model.

In connection with future model re-estimations, we emphasize the following points: To estimate the model, data must contain measurements of travel time and traffic flow before, during, and after the peak periods for each day in the analysis period. This is necessary to identify when congestion sets in and to analyse the dynamic process that leads to its dissolution. Travel times measured at segment or link level are preferred to travel times measured at point level (we apply the latter in our analysis from lack of better options). GPS or Bluetooth data of travel times are most likely much more accurate than loop detector data and should be applied if they have sufficient data coverage. To compute TTV at a given time of day it is necessary with repeated observations of travel times at this time, on a given day and over many days.

Following the discussion of spillback effects in the preceding section, we emphasize that it is necessary to observe traffic conditions not only on the road links of interest, but also on adjacent (in particular downstream) road links, preferably including entry and exit ramps and divergences.

An important, though somewhat obvious, remark is that we can only estimate road-specific parameters for a road, if the road in question is sufficiently congested. If congestion does not occur with the current traffic demand, it is impossible to identify the demand levels at which congestion would occur.¹¹

Finally, a lesson from the project is that the data work connected to filtering out data points affected by road works and short term maintenance works turned out to be extensive. In connection with future data collection we recommend that such incidents are systematically logged and registered in both Mastra and Hastrid, or in a separate database which can easily be merged with these.

¹¹ Note that this is not the same as to say that the model cannot be applied to forecast TTV on roads without sufficient congestion: We can use the model for this, but one has to apply parameters that are estimated on data for a different (similar) road.

References

The Danish Road Directorate (2010a): Vejregel for kapacitet og serviceniveau, version september 2010.

The Danish Road Directorate (2010b): Brugervejledning til Hastrid, version juli 2010.

de Jong, G.C., Bliemer, M.C.J. (2015): On including travel time reliability of road traffic in appraisal, *Transportation Research Part A*, vol 73, pp. 80-95.

de Palma, A., Fosgerau, M. (2011): Dynamic Traffic Modeling, chapter 9 in *A Handbook of Transport Economics*. Edward Elgar Publishing, UK.

DTU Transport: Travel time variability - Definition and valuation, Report 1-2008.

Fosgerau, M., Small, K. (2012): Marginal Congestion Cost on a Dynamic Expressway Network, *Journal of Transport Economics and Policy*, vol. 46(3), pp. 431-450.

Overgård Hansen, C., Filges, D., Sørensen, H., Brun, B. (2014): Modellering af kø-opstuvning i en statisk rutevalgmodel, Trafikdage på Aalborg Universitet 2014.

Tetraplan (2013): Flaskehalse – Dokumentation af fase 3 og database.

Wooldridge, J. M. (2002): *Econometric analysis of cross section and panel data*, MIT Press.

Appendix A Tables

Table 13: Estimation results for breakdown model.

Sample	Max log likelihood	#Obs	β_0	robust std.err. (β_0)	β_1	robust std.err. (β_1)
Links 1,2,8	-554.68	5585	-13.689***	0.714	0.399***	0.023
Links 1,2,3,8	-893.95	7160	-11.350***	0.455	0.316***	0.015
Link 1	-126.34	1653	-21.736***	2.399	0.647***	0.075
Link 2	-179.99	1512	-19.041***	1.673	0.558***	0.052
Link 3	-316.84	1575	-8.503***	0.614	0.218***	0.020
Link 4	-95.12	4569	-11.059***	1.259	0.222***	0.045
Link 5			--- Did not converge ---			
Link 8	-209.73	2420	-11.936***	1.235	0.367***	0.047

*** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 14: Estimation results for recovery model.

Sample	Model	κ	Max log likelihood	#Obs	γ_0	robust std.err. (γ_0)	γ_1	robust std.err. (γ_1)	γ_2	robust std.err. (γ_2)
Links 1,2,8	eq. (9)	20	-662.99	1640	0.875	0.532	0.167	0.724	0.063**	0.027
	eq. (9)	21	-663.06	1640	1.792***	0.483	-0.325	0.778	0.081***	0.029
	eq. (9)	22	-661.63	1640	2.079***	0.401	-0.767	0.833	0.097***	0.032
	eq. (9)	23	-660.50	1640	1.938***	0.252	-1.414	0.964	0.120***	0.036
	eq. (10)	20	-663.26	1640	0.875	0.532	-3.225	2.300	1.544**	0.704
	eq. (10)	21	-663.23	1640	1.792***	0.483	-4.985**	2.497	2.077***	0.764
	eq. (10)	22	-661.70	1640	2.079***	0.401	-6.545**	2.694	2.549***	0.824
	eq. (10)	23	-660.40	1640	1.938***	0.252	-8.907***	3.154	3.261***	0.963
Links 1,2,3,8	eq. (9)	20	-1023.14	2598	1.030**	0.521	1.216**	0.543	0.024	0.020
	eq. (9)	21	-1023.71	2598	1.887***	0.480	0.985*	0.566	0.032	0.020
	eq. (9)	22	-1022.96	2598	2.148***	0.399	0.805	0.589	0.038*	0.021
	eq. (9)	23	-1022.90	2598	1.977***	0.251	0.663	0.640	0.043*	0.023
	eq. (10)	20	-1023.12	2598	1.030**	0.521	-0.321	1.756	0.662	0.530
	eq. (10)	21	-1023.59	2598	1.887***	0.480	-1.178	1.848	0.917	0.558
	eq. (10)	22	-1022.74	2598	2.148***	0.399	-1.850	1.938	1.118*	0.585
	eq. (10)	23	-1022.61	2598	1.977***	0.251	-2.429	2.126	1.290**	0.641
Link 1	eq. (9)	20	-206.89	532			---	not identified	---	
	eq. (9)	21	-205.81	532			---	not identified	---	
	eq. (9)	22	-205.63	532	2.944***	1.026	-0.166	1.586	0.079	0.062
	eq. (9)	23	-206.20	532	2.007***	0.355	-0.810	1.956	0.103	0.076
	eq. (10)	20	-206.94	532			---	not identified	---	
	eq. (10)	21	-205.86	532			---	not identified	---	
	eq. (10)	22	-205.67	532	2.944***	1.026	-4.547	5.140	1.979	1.589
	eq. (10)	23	-206.23	532	2.007***	0.355	-6.758	6.409	2.652	1.972
Link 2	eq. (9)	20	-295.95	805			---	not identified	---	
	eq. (9)	21	-295.95	805			---	not identified	---	
	eq. (9)	22	-295.52	805			---	not identified	---	
	eq. (9)	23	-295.06	805	3.045***	1.024	0.889	1.247	0.039	0.045
	eq. (10)	20	-295.98	805			---	not identified	---	
	eq. (10)	21	-295.98	805			---	not identified	---	
	eq. (10)	22	-295.55	805			---	not identified	---	
	eq. (10)	23	-295.09	805	3.045***	1.024	-1.446	4.083	1.030	1.232
Link 3	eq. (9)	20	-356.26	958			---	not identified	---	
	eq. (9)	21	-356.18	958			---	not identified	---	
	eq. (9)	22	-356.09	958			---	not identified	---	
	eq. (9)	23	-356.09	958			---	not identified	---	
	eq. (10)	20	-356.26	958			---	not identified	---	
	eq. (10)	21	-356.18	958			---	not identified	---	
	eq. (10)	22	-356.09	958			---	not identified	---	
	eq. (10)	23	-355.99	958			---	not identified	---	
Link 8	eq. (9)	20	-154.11	303	0.788	0.539	-1.982	1.624	0.131**	0.063
	eq. (9)	21	-150.78	303	1.526***	0.493	-5.938***	2.264	0.282***	0.088
	eq. (9)	22	-145.79	303	1.705***	0.444	-10.697***	3.067	0.465***	0.120
	eq. (9)	23	-145.79	303	1.531***	0.390	-12.579***	3.470	0.537***	0.136
	eq. (10)	20	-154.44	303	0.788	0.539	-7.923	4.888	2.867*	1.510
	eq. (10)	21	-151.08	303	1.526***	0.493	-20.844***	7.021	6.831***	2.166

	eq. (10)	22	-145.78	303	1.705***	0.444	-37.010***	9.717	11.790***	2.999
	eq. (10)	23	-144.94	303	1.531***	0.390	-43.762***	11.002	13.861***	3.396

*** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 15: Estimation results for recovery model with γ_0 fixed to zero.

Sample	Model	κ	Max log likelihood	#Obs	γ_1	robust std.err. (γ_1)	γ_2	robust std.err. (γ_2)
Link1	eq. (9)	20	-206.89	532	0.795	1.398	0.043	0.055
	eq. (9)	21	-209.97	532	0.336	1.481	0.060	0.058
	eq. (9)	22	-215.52	532	-0.166	1.586	0.079	0.062
	eq. (9)	23	-231.23	532	-0.804	1.956	0.103	0.076
	eq. (10)	20	-206.94	532	-1.278	4.473	0.979	1.384
	eq. (10)	21	-210.02	532	-2.862	4.776	1.465	1.478
	eq. (10)	22	-215.56	532	-4.547	5.140	1.979	1.589
	eq. (10)	23	-231.26	532	-6.759	6.409	2.652	1.972
Link 2	eq. (9)	20	-296.64	805	1.464	1.142	0.019	0.041
	eq. (9)	21	-296.64	805	1.464	1.142	0.019	0.041
	eq. (9)	22	-298.29	805	1.338	1.162	0.023	0.042
	eq. (9)	23	-306.24	805	0.889	1.247	0.039	0.045
	eq. (10)	20	-296.67	805	0.548	3.705	0.435	1.119
	eq. (10)	21	-296.67	805	0.548	3.705	0.435	1.119
	eq. (10)	22	-298.32	805	0.099	3.781	0.569	1.142
	eq. (10)	23	-306.27	805	-1.446	4.083	1.030	1.232
Link 3	eq. (9)	20	-357.65	958	3.438***	1.056	-0.050	0.035
	eq. (9)	21	-358.26	958	3.409***	1.063	-0.049	0.035
	eq. (9)	22	-358.86	958	3.382***	1.068	-0.048	0.035
	eq. (9)	23	-358.86	958	3.382***	1.068	-0.048	0.035
	eq. (10)	20	-357.64	958	6.929**	3.498	-1.465	1.029
	eq. (10)	21	-358.25	958	6.825*	3.527	-1.435	1.038
	eq. (10)	22	-358.86	958	6.729*	3.549	-1.407	1.044
	eq. (10)	23	-359.46	958	6.638*	3.566	-1.380	1.049

*** denotes significance at the 1% level, ** at the 5% level, and * at the 10% level.

Acknowledgements

Jan Holm and Jakob Elbek, The Danish Road Directorate.

We are grateful to Jan and Jakob for all their help with the Mastra and Hastrid data.

Jens Foller, The Danish Road Directorate, and **Lars Olsen**, The Danish Ministry of Transport. Jens and Lars has followed the project from beginning to end, and kindly shared their expertise regarding both practical and theoretical matters.

DTU Transport forsker og underviser i trafik og transportplanlægning. Institutet rådgiver myndighederne inden for infrastruktur, samfundsøkonomi, transportpolitik og trafiksikkerhed. DTU Transport samarbejder tillige med erhvervslivet om grøn logistik, behovsstyret kollektiv trafik, brugerbetaling og design af bæredygtige transportnetværk.

DTU Transport
Institut for Transport
Danmarks Tekniske Universitet

Bygningstorvet 116B
2800 Kgs. Lyngby
Tlf. 45 25 65 00
Fax 45 93 65 33

www.transport.dtu.dk