

Revenue Management bei nachfrage- und anbieterseitiger Substitution: Methoden der approximativen dynamischen Optimierung

Sebastian Koch

Angaben zur Veröffentlichung / Publication details:

Koch, Sebastian. 2017. "Revenue Management bei nachfrage- und anbieterseitiger Substitution: Methoden der approximativen dynamischen Optimierung." Augsburg: Universität Augsburg.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under the following conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publizieren>



Revenue Management bei nachfrage- und anbieterseitiger Substitution

Methoden der approximativen dynamischen Optimierung

Kumulative Dissertation
der Wirtschaftswissenschaftlichen Fakultät
der Universität Augsburg
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
(Dr. rer. pol.)

vorgelegt von

M.Sc. with honors Sebastian Koch

Erstgutachter:	Prof. Dr. Robert Klein
Zweitgutachter:	Prof. Dr. Florian Jaehn
Vorsitzender der mündlichen Prüfung:	Prof. Dr. Axel Tuma
Datum der mündlichen Prüfung:	19.12.2016

Verzeichnis der Beiträge

Diese Dissertation umfasst die folgenden veröffentlichten bzw. zur Veröffentlichung angenommenen Beiträge. Die jeweils angegebenen Kategorien beziehen sich auf das Zeitschriftenranking JOURQUAL 3 des Verbands der Hochschullehrer für Betriebswirtschaft e. V. (VHB). Die Sortierung entspricht der Reihenfolge des Abdrucks im vorliegenden Manuskript.

B1 Beitrag der Kategorie B

Koch S (2017) Least squares approximate policy iteration for learning bid prices in choice-based revenue management. *Computers & Operations Research* 77(2017): 240–253.

B2 Beitrag der Kategorie A

Koch S, Gönsch J, Steinhardt C (2017) Dynamic programming decomposition for choice-based revenue management with flexible products. Erscheint in: *Transportation Science*.

B3 Beitrag der Kategorie B

Gönsch J, Koch S, Steinhardt C (2014) Revenue management with flexible products – The value of flexibility and its incorporation into DLP-based approaches. *International Journal of Production Economics* 153(2014): 280–294.

B4 Beitrag der Kategorie B

Gönsch J, Koch S, Steinhardt C (2013) An EMSR-based approach for revenue management with integrated upgrade decisions. *Computers & Operations Research* 40(10): 2532–2542.

B5 Beitrag der Kategorie C

Koch S, Gönsch J, Hassler M, Klein R (2016) Practical decision rules for risk-averse revenue management using simulation-based optimization. *Journal of Revenue and Pricing Management* 15(6): 468–487.

Inhaltsverzeichnis

I	Einleitung	1
I.1	Einführung in Revenue Management und (approximative) dynamische Optimierung.....	2
I.2	Revenue Management bei nachfrageseitiger Substitution	6
I.3	Revenue Management bei anbieterseitiger Substitution	8
I.4	Revenue Management bei Risikoaversion	12
II	Beiträge.....	17
	Beitrag B1: Least squares approximate policy iteration for learning bid prices in choice-based revenue management	18
	Beitrag B2: Dynamic programming decomposition for choice-based revenue management with flexible products	19
	Beitrag B3: Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches.....	78
	Beitrag B4: An EMSR-based approach for revenue management with integrated upgrade decisions	79
	Beitrag B5: Practical decision rules for risk-averse revenue management using simulation-based optimization	80
III	Fazit und Ausblick.....	111

I Einleitung

Im Zuge der Deregulierung des US-amerikanischen Markts für Passagierluftfahrt Ende der Siebzigerjahre drängten zahlreiche neue Wettbewerber mit attraktiven Preisen auf den Markt. Diese konnten ihre vergleichsweise günstigen Preise vor allem aufgrund geringerer Kosten anbieten, bspw. in Folge eines weitgehenden Verzichts auf Serviceleistungen. Die etablierten Fluggesellschaften standen vor der Herausforderung, eine Abwanderung insbesondere von preissensitiven Privatreisenden zur Konkurrenz zu verhindern. Jedoch konnten sie ihre Preise aufgrund der strukturell bedingt höheren Kosten nicht generell auf das Niveau der Konkurrenz absenken. So griff American Airlines auf das Instrument der Preisdifferenzierung zurück und führte den speziell an Privatreisende gerichteten, wettbewerbsfähigen Tarif „Ultimate Super Saver Fare“ ein, der an die Bedingung einer Vorausbuchungsfrist geknüpft war. Nun konkurrierten unterschiedlich bepreiste Produkte um dieselbe Ressource. Damit stand American Airlines vor der Fragestellung, wie viele der günstigen Tickets höchstens verkauft werden durften, ohne die meist später eintreffende, aber lukrativere Nachfrage von Geschäftsreisenden zu verdrängen. Diese offene Fragestellung gilt als Geburtsstunde des Revenue Managements, das heute neben der Passagierluftfahrt in den verschiedensten Industriezweigen erfolgreich angewendet wird, bspw. in der Automobilvermietung (vgl. Steinhardt and Gönsch (2012)), der Medienwirtschaft (vgl. Kimms und Müller-Bungart (2007)) oder dem Attended Home Delivery (vgl. Agatz et al. (2013) und Yang et al. (2016)).

Im Allgemeinen befasst sich Revenue Management mit der optimalen Gestaltung zeitlich begrenzter Verkaufsprozesse kapazitätsbeschränkter Ressourcen. Die resultierenden Entscheidungsprobleme werden üblicherweise als stochastische, dynamische Optimierungsprobleme (engl. Dynamic Program, DP) mit endlichem Zeithorizont modelliert. Da praxisrelevante Instanzen dieser Art von Problemstellung meist nicht exakt gelöst werden können, muss auf approximative Lösungsmethoden zurückgegriffen werden. Diese sind üblicherweise auf das vorliegende Entscheidungsproblem angepasste Weiterentwicklungen existierender Methoden.

Im vergangenen Jahrzehnt waren im Revenue Management zwei wegweisende Trends beobachtbar: die Berücksichtigung nachfrageseitiger Substitution durch Kundenwahlverhalten sowie anbieterseitiger Substitution, bspw. durch flexible Produkte und Upgrades. Ein dritter Aspekt, der vor allem aus dem Transfer von Revenue Management in andere Anwendungsfelder motiviert ist, besteht in der Betrachtung eines risikoaversen Entscheiders.

Gegenstand der vorliegenden kumulativen Dissertationsschrift ist die (approximative) Lösung von Entscheidungsproblemen diese drei Trends betreffend. Vor diesem Hintergrund erfolgt in Kapitel I.1 zunächst eine Einführung in das Forschungsgebiet des klassischen Revenue Managements inklusive entsprechender (approximativer) Lösungsmethoden. In den Kapitel I.2 bis I.4 werden die Trends näher vorgestellt und die in dieser Dissertationsschrift vorgestellten Beiträge inhaltlich verortet.

I.1 Einführung in Revenue Management und (approximative) dynamische Optimierung

Die vorliegende Dissertationsschrift versteht den Begriff des Revenue Managements im engeren Sinne einer klassischen Kapazitätssteuerung, wie es bspw. der folgenden Definition von Klein (2001) entspricht:

„Revenue Management umfasst eine Reihe von quantitativen Methoden zur Entscheidung über Annahme oder Ablehnung unsicherer, zeitlich verteilt eintreffender Nachfrage unterschiedlicher Wertigkeit. Dabei wird das Ziel verfolgt, die in einem begrenzten Zeitraum verfügbare, unflexible Kapazität möglichst effizient zu nutzen.“

Aus Anbietersicht lassen sich die meisten Problemstellungen im Revenue Management im Wesentlichen wie folgt charakterisieren (vgl. Talluri und van Ryzin (2004b), Kap. 1.3.3 oder Klein und Steinhardt (2008), Kap. 1.2.2):

- Die abzusetzenden Produkte entstehen im Rahmen einer segmentorientierten Preisdifferenzierung zweiten oder dritten Grades (vgl. Phillips (2005), Kap. 4 sowie Klein und Steinhardt (2008), Kap. 2 zum Begriff und zu Ausprägungen der Preis-

differenzierung). Die Produkte sind somit in ihrer Kernleistung identisch, aber unterschiedlich bepreist.

- Die potentiellen Kunden besitzen heterogene Präferenzen bzgl. des Erwerbs der Produkte und unterschiedliche Zahlungsbereitschaften.
- Der Absatz der Produkte erfolgt über einen zeitlich begrenzten Verkaufszeitraum, der meist als Buchungshorizont bezeichnet wird. Die eigentliche Leistungserstellung hingegen findet nach dem Buchungshorizont statt.
- Die Leistungserstellung erfordert von den Produkten gemeinsam genutzte Ressourcen. Die Ressourcen sind kapazitätsbeschränkt und nicht lagerfähig. Betrachtet man lediglich eine Ressource, bspw. in der Passagierluftfahrt einen Flugabschnitt (engl. Leg), spricht man vom Single-Leg-Fall, andernfalls vom Netzwerk-Fall.

In einem derartigen Umfeld ist es nun die Aufgabe des Revenue Managements im Sinne einer Kapazitätssteuerung, dynamisch über die Verfügbarkeit der Produkte im Buchungshorizont zu entscheiden, so dass ein zuvor festgelegtes Zielkriterium optimiert wird. Für gewöhnlich handelt es sich dabei um den Erwartungswert des über den Buchungshorizont erzielten Erlöses. Das Ziel der Erlösmaximierung ist dabei aufgrund der Kostenstruktur in traditionellen Anwendungsfeldern wie der Passagierluftfahrt gerechtfertigt, die durch hohe Fix- und geringe variablen Kosten geprägt ist. Die Optimierung des Erwartungswerts basiert auf der Annahme eines risikoneutralen Entscheiders, wie es in vielen Anwendungsfeldern aufgrund einer großen Anzahl ähnlicher Buchungsprozesse angebracht ist.

Die Modellierung dieser Entscheidungsprobleme als DP hat sich mittlerweile als Standardansatz etabliert. Da im Rahmen dieser Einführung nur die Grundprinzipien der dynamischen Optimierung erläutert werden können, seien für einen ausführlicheren Einstieg in die Thematik Domschke et al. (2015), Kap. 7 und Nickel et al. (2011), Kap. 8 empfohlen. Allgemein besteht ein DP aus einer Menge von Perioden, Systemzuständen, Aktionen, einperiodigen Zielfunktionsbeiträgen und Übergangswahrscheinlichkeiten (vgl. Puterman (2005), Kap. 2.1). In jeder Periode ist im gegenwärtigen Zustand des Systems eine Entscheidung zu treffen, d. h. eine Aktion aus der Menge der zulässigen

Aktionen zu wählen. Eine gewählte Aktion führt in Abhängigkeit des gegenwärtigen Zustands und der Übergangswahrscheinlichkeiten zu einem Zielfunktionsbeitrag sowie in einen Folgezustand. Die Wahrscheinlichkeit, in einen bestimmten Folgezustand zu gelangen, hängt dabei nicht von vergangenen Zuständen oder Entscheidungen ab – eine Eigenschaft, die häufig als Markov-Eigenschaft bezeichnet wird (vgl. Hillier und Lieberman (2010), Kap. 16.2). Eine Abfolge von Aktionen bzw. eine Regel zur Bestimmung einer solchen Abfolge wird als Politik bezeichnet (vgl. Puterman (2005), Kap. 2.1.5). Ziel ist die Bestimmung einer optimalen Politik, d. h. einer Politik, die für jeden Zustand in jeder Periode eine optimale Aktion hinsichtlich des zuvor festgelegten Zielkriteriums angibt. Die Markov-Eigenschaft ermöglicht die Modellierung des Entscheidungsproblems nur in Abhängigkeit von gegenwärtigem Zustand und gegenwärtiger Periode. Dazu wird eine sog. Wertfunktion formuliert, die den optimalen Zielfunktionswert in diesem Zustand ab dieser Periode angibt. Die Lösung erfolgt mittels Rückwärtsrechnung, indem für jeden Zustand jeder Periode die Wertfunktion berechnet wird.

Zur Modellierung von Entscheidungsproblemen des Revenue Managements als DP wird zunächst der Buchungshorizont in eine endliche Anzahl hinreichend kleiner Perioden diskretisiert, so dass in jeder Periode maximal ein Kunde eintrifft (vgl. Subramanian et al. (1999) zur Diskretisierung des Buchungshorizonts). Die Nachfrage wird somit disaggregiert, d. h. auf kundenindividueller Ebene, modelliert. Der Anbieter entscheidet in jeder Periode über die anzubietenden Produkte. Die Entscheidung, ob ein bestimmtes Produkt angeboten werden soll, hängt dabei maßgeblich von den Opportunitätskosten des Produkts ab, d. h. dem entgangenen erwarteten Erlös im Falle eines Verkaufs des Produkts (vgl. Klein und Steinhardt (2008), Kap. 3.1.2 zum Begriff der Opportunitätskosten im Revenue Management). Von den angebotenen Produkten fragt ein Kunde maximal eines nach. Der in einer Periode erzielte Erlös sowie der Zustandsübergang von einer Periode zur Folgeperiode erfolgt folglich in Abhängigkeit der angebotenen Produkte sowie der Kaufwahrscheinlichkeiten. Dabei wird – zumindest im Standardfall – der Zustand des Systems durch die verfügbare Restkapazität der zur Leistungserstellung erforderlichen Ressourcen beschrieben. Bei Verkauf eines bestimmten Produkts

verringert sich die Restkapazität um den entsprechenden Kapazitätsbedarf. Die Opportunitätskosten eines Produkts spiegeln exakt den Wert dieses Kapazitätsverbrauchs wider.

Wie eingangs beschrieben, ist die exakte Lösung derartiger DPs allerdings nur für kleine Probleminstanzen möglich. Dies liegt vor allem am sog. „curse of dimensionality“, der verhindert, dass die Wertfunktion für den vollständigen Zustandsraum, d. h. für sämtliche denkbaren Zustände jeder Periode, berechnet werden kann (vgl. Powell (2011), Kap. 1.2 zum „curse of dimensionality“). Um dennoch sinnvolle Entscheidungen über die anzubietenden Produkte treffen zu können, existiert eine Vielzahl approximativer Lösungsmethoden, die sich im Wesentlichen wie folgt klassifizieren lassen:

- Kapazitätssteuerung mittels Buchungslimits (bzw. Schutzlimits): Die Grundidee dieser Revenue Management-exklusiven Methode besagt, dass ein Produkt solange angeboten wird, bis eine vorher festgelegte maximale Absatzmenge, das sog. Buchungslimit, erschöpft ist. Die Ermittlung erfolgt in der Regel mit Hilfe eines einfach lösbaren deterministischen Ersatzmodells wie des bekannten Erwartungswertmodells (engl. Deterministic linear program, DLP) oder den von Belobaba (1987, 1989, 1992) vorgeschlagenen Expected Marginal Seat Revenue (EMSR) Verfahren EMSR-a und EMSR-b. Diese Art der Kapazitätssteuerung kann alternativ über sog. Schutzlimits ausgedrückt werden, die die vor dem Zugriff durch ein Produkt zu schützenden Kapazitätseinheiten beschreiben. Buchungs- und Schutzlimits können einfach ineinander überführt werden.
- Kapazitätssteuerung mittels Bid-Preisen: Diese Methode basiert auf dem Vergleich des Erlöses eines Produkts mit einer Preisuntergrenze. Diese Entscheidungsregel ist aus dem DP des klassischen Revenue Managements ohne nachfrageseitige Substitution abgeleitet. In diesem vergleichsweise einfach zu handhabendem Anwendungsfall lohnt es sich, ein Produkt anzubieten, falls der Erlös des Produkts die Opportunitätskosten übersteigt (vgl. Talluri und van Ryzin (2004b), Kap. 3.2.1). Die Preisuntergrenze stellt folglich eine Approximation der eigentlichen Opportunitätskosten des Produkts dar. Die Preisuntergrenze wird mit Hilfe sog. Bid-Preise berechnet. Ein

Bid-Preis bezeichnet dabei die monetäre Bewertung einer Kapazitätseinheit einer Ressource, so dass sich die Preisuntergrenze aus der Summe der mit dem Kapazitätsbedarf gewichteten Bid-Preise der benötigten Ressourcen ergibt.

- Approximation der Wertfunktion: Dieses (approximative) Lösungsprinzip für DPs basiert auf der Verwendung einer zu lernenden Approximation der ursprünglichen Wertfunktion. Es kommt häufig auch in anderen Anwendungsfeldern der dynamischen Optimierung zum Einsatz. Der Einsatz von Wertfunktions-Approximationen wird dabei oft als approximative dynamische Optimierung (engl. approximate dynamic programming, ADP) im engeren Sinne verstanden. Anzumerken ist jedoch, dass die Abgrenzung des Begriffs ADP zu anderen approximativen Lösungsmethoden oft nicht eindeutig ist (vgl. Powell (2011), Kap. 6).

Die vorgestellte Klassifizierung der Methoden ist häufig nicht trennscharf. Beispielsweise erfolgt die Formulierung der meisten Wertfunktions-Approximationen über Bid-Preise. Ein guter, allgemeiner Überblick insbesondere über die Kapazitätssteuerung mittels Buchungslimits und Bid-Preisen sowie über gängige Ersatzmodelle findet sich in Talluri und van Ryzin (2004b), Kap. 2 und 3 sowie in Klein und Steinhardt (2008), Kap. 3. Die Eignung von Wertfunktions-Approximationen zur Kapazitätssteuerung wird vor allem in neueren Veröffentlichungen untersucht, bspw. in Adelman (2007) oder Meissner und Strauss (2012).

I.2 Revenue Management bei nachfrageseitiger Substitution

In früheren Anwendungen des Revenue Managements wurde vereinfachend angenommen, dass die Nachfrage nach einzelnen Produkten stochastisch unabhängig ist, d. h. die Produkte aus Kundensicht keine Substitute darstellen. Diese sog. Independent Demand-Annahme erleichterte das in jeder Periode des DPs zu lösende Entscheidungsproblem erheblich, da für jedes Produkt separat entschieden werden konnte, ob sich ein Anbieten lohnt oder nicht (vgl. Kap. I.1). Ausgehend von Gallego et al. (2004) sowie Talluri und van Ryzin (2004a) hat sich jedoch mittlerweile die Berücksichtigung von nachfrageseitiger Substitution aufgrund von Kundenwahlverhalten etabliert. In diesem Fall resultiert

das in jeder Periode zu lösende Entscheidungsproblem in einem (kombinatorischen) Sortimentsoptimierungsproblem, dessen Komplexität stark vom unterstellten Kundenwahlverhalten abhängt. Nicht nur im Revenue Management am weitesten verbreitet ist dabei das multinomiale Logit-Modell (MNL-Modell). Die Entwicklung des Modells wurde maßgeblich von Nobelpreisträger Daniel McFadden geprägt (vgl. McFadden (2001)). Das MNL-Modell gilt als der am einfachsten zu handhabende Vertreter der sog. Discrete Choice Modelle (vgl. Train (2009), Kap. 3). Während Kök et al. (2009) einen Überblick über Sortimentsoptimierung geben, analysieren unter anderem Liu und van Ryzin (2008) sowie Miranda Bront et al. (2009) Varianten des sich bei Annahme eines MNL-Modells ergebenden Sortimentsoptimierungsproblems.

Beitrag B1 „Least squares approximate policy iteration for learning bid prices in choice-based revenue management“ greift eine Forschungslücke im Revenue Management bei nachfrageseitiger Substitution auf. Zwar existieren bereits einige Methoden zum Lernen von Wertfunktions-Approximationen, diese basieren jedoch fast ausschließlich auf analytisch motivierten Optimierungsansätzen. Das in den letzten Jahren vor allem in anderen Anwendungsfeldern populär gewordene simulationsbasierte ADP, über das Powell (2011) einen guten Überblick gibt, findet bisher hingegen wenig Beachtung. Die Grundidee dieser Verfahren ist es, beginnend im Startzustand bis zur letzten Periode abwechselnd eine Aktion auswählen und den Folgezustand des Systems zu simulieren. Die Auswahl der Aktion erfolgt dabei unter Berücksichtigung der aktuellen Approximation. Einen solchen Durchlauf bezeichnet man als Simulationspfad. Die Verbesserung der Approximation geschieht im Nachgang unter Verwendung der auf dem Simulationspfad erzielten Erlöse, die Schätzer für die Approximation in den im Simulationspfad besuchten Zuständen sind. Im Beitrag werden lineare Approximationen untersucht, so dass zur Verbesserung der Approximationen lediglich Kleinste-Quadrate-Probleme gelöst werden müssen (vgl. Brooks (2014), insbesondere Kap. 2, zur Methode der kleinsten Quadrate). Um im Gegensatz zu den vorherrschenden analytischen Ansätzen kein bestimmtes Kundenwahlverhalten vorauszusetzen, wird zur Sortimentsoptimierung im Rahmen der Simulation ein von Prokopyev (2005) und Miranda Bront et al.

(2009) vorgeschlagener Greedy-Algorithmus herangezogen. Umfangreiche Tests auf Grundlage von Standard-Probleminstanzen der aktuellen Revenue Management-Literatur zeigen, dass der resultierende Ansatz mehr als nur kompetitive Ergebnisse im Vergleich zu bestehenden Methoden liefert.

1.3 Revenue Management bei anbieterseitiger Substitution

Ein zweiter wichtiger Trend in der Literatur zum Revenue Management ist die Berücksichtigung anbieterseitiger Substitutionsmöglichkeiten. Dieser Trend wurde maßgeblich vorangetrieben durch den Aufstieg des Internets als Verkaufskanal und damit einhergehend der Etablierung neuartiger Produkte – so z. B. das Anbieten unvollständig spezifizierter Produkte, deren genaue Ausgestaltung ein Kunde erst nach dem Kauf erfährt. Die vorliegende Dissertationsschrift befasst sich in diesem Zusammenhang mit flexiblen Produkten sowie Upgrades, die als zwei der bekanntesten Beispiele solcher innovativer Produkte gelten.

Ein flexibles Produkt bezeichnet ein Bündel mehrerer substituierbarer Alternativen, von denen der Verkäufer dem Käufer eine Alternative an einem zuvor vereinbarten Zeitpunkt nach dem Kauf zuweist (vgl. Gallego und Phillips (2004)). Flexible Produkte bieten dem Anbieter zwei entscheidende Vorteile:

- Nachfrageinduktion: Von Kundenseite wird ein flexibles Produkt aufgrund der Unsicherheit der endgültigen Ausgestaltung als inferior im Vergleich zu herkömmlichen, regulären Produkten empfunden. Damit ermöglichen flexible Produkte die Erschließung neuer Kundengruppen mit geringer Zahlungsbereitschaft, ohne zu viel hochwertige Nachfrage nach regulären Produkten zu kannibalisieren (vgl. Jerath et al. (2010)).
- Flexibilität bzgl. der Ressourcenbelegung sowie dem Zeitpunkt der Zuweisung: Aus der Sicht des Revenue Managements für den Anbieter besonders interessant ist nicht nur die Möglichkeit, mehrere Zuweisungsalternativen zur Auswahl zu haben, sondern vor allem, die endgültige Alternative erst zu einem späteren Zeitpunkt als dem Verkaufszeitpunkt festzulegen. Die Zuweisung kann bspw. zu einem Zeitpunkt

erfolgen, an dem bereits ein Großteil der Nachfrage eingetroffen ist und die verbleibende Unsicherheit hinsichtlich der noch eintreffenden Nachfrage daher erheblich geringer ist als zum Zeitpunkt des Verkaufs. Flexible Produkte tragen somit zur Verbesserung der Kapazitätsauslastung bei, insbesondere im Falle einer schwer zu prognostizierenden Nachfrage (vgl. Petrick et al. (2012)). Nachfolgend wird die Möglichkeit der Verschiebung der Zuweisungsentscheidung zu einem späteren Zeitpunkt als zeitliche Flexibilität bezeichnet.

Flexible Produkte finden sich in der Praxis vor allem in der Tourismusbranche. Beispielsweise bietet Aida Cruises (www.aida.com) gegenwärtig drei flexible Produkte an: „JUST AIDA“, „JUST AIDA First Minute“ und „AIDA VARIO“. Alle drei Produkte bieten dem Anbieter den Vorteil der zeitlichen Flexibilität, verbergen zum Verkaufszeitpunkt aber unterschiedliche Eigenschaften vor dem Kunden. Sowohl bei „JUST AIDA“ als auch „JUST AIDA First Minute“ wird je nach Kundenpräferenz entweder der genaue Reiseternin oder die Reiseroute verschwiegen. Bei „AIDA VARIO“ hat der Anbieter lediglich einen gewissen Spielraum bei der Kabinenauswahl.

Von Upgrades spricht man, wenn ein Kunde zum Zeitpunkt der Leistungserstellung unentgeltlich ein höherwertigeres Produkt als das ursprünglich gekaufte erhält. Hierbei geht man üblicherweise davon aus, dass ein Kunde ein solches Upgrade immer akzeptiert. Die Vergabe von Upgrades ist vor allem in der Passagierluftfahrt (vgl. Gönsch und Steinhardt (2015)) und der Automobilvermietung (vgl. Fink und Reiners (2006)) verbreitet. So erhalten Kunden bei Kapazitätsengpässen eine höherwertigere Beförderungsklasse (bspw. Business statt Economy) bzw. Fahrzeugkategorie (bspw. Intermediate statt Economy). Ähnlich wie flexible Produkte bieten Upgrades dem Anbieter eine erweiterte Flexibilität in der Kapazitätssteuerung.

Über flexible Produkte und Upgrades hinaus gibt es eine Reihe von verwandten Konzepten, bei denen die Substitutionsmöglichkeit implizit vorgesehen und für den Kunden nicht sichtbar oder nicht von Bedeutung ist. Beispielhaft zu nennen ist der Transport von Luftfracht, bei dem die endgültige Routenführung dem Transportdienstleister überlassen ist (vgl. Bartodziej et al. (2007)). Weitere Beispiele finden sich in der Medien-

wirtschaft (vgl. Müller-Bungart (2007)) sowie der Auftragsfertigung (vgl. Guhlich et al. (2015)). In der Literatur zum Revenue Management spricht man oft auch in solchen Fällen von flexiblen Produkten, da die Modellierung analog erfolgt und die Betrachtung der erweiterten Flexibilität in der Kapazitätssteuerung im Vordergrund steht.

In der vorliegenden Dissertationsschrift wird der einfacheren Darstellung wegen angenommen, dass die Zuweisungsentscheidung zu einer endgültigen Alternative sowohl bei flexiblen Produkten als auch bei Upgrades erst zum Zeitpunkt der Leistungserstellung erfolgt.

Obwohl insbesondere Upgrades in der Praxis seit langem verbreitet sind, werden die Auswirkungen von anbieterseitigen Substitutionsmöglichkeiten auf die Methoden des Revenue Managements erst seit knapp über einem Jahrzehnt erforscht – beginnend mit den richtungsweisenden wissenschaftlichen Artikeln von Gallego und Phillips (2004) sowie Gallego et al. (2004). Um flexible Produkte und Upgrades in DPs zu integrieren, reicht es nicht mehr aus, dass ein Zustand lediglich durch die verfügbare Restkapazität der Ressourcen beschrieben wird. Stattdessen müssen die bereits verkauften Einheiten (Reservierungen) der Produkte in den Zustandsraum integriert werden. Die fehlende eindeutige Beziehung zu einer konkreten Alternative und damit zu einem eindeutigen Ressourcenverbrauch bringt vor allem zwei Probleme mit sich, die einen Einsatz der Standardmethoden (vgl. Kapitel I.1) zunächst verhindern. Zum einen gestaltet sich die Bestimmung von approximativen Opportunitätskosten der Produkte schwierig, da dafür im Gegensatz zu regulären Produkten keine Bid-Preise verwendet werden können. Zum anderen muss zu jedem Zeitpunkt des Buchungshorizonts sichergestellt sein, dass die Kapazität zur Erfüllung sämtlicher Reservierungen ausreicht. Im zugehörigen DP geschieht dies über ein inhärentes Zulässigkeitsproblem. In bisherigen (heuristischen) Ansätzen wurden diese beiden Probleme meistens dadurch gelöst, dass die Reservierungen bereits im Verkaufszeitpunkt entweder endgültig oder zumindest temporär bestimmten Alternativen zugeordnet werden (vgl. Petrick et al. (2010), Petrick et al. (2012) sowie Steinhardt und Gönsch (2012)), was mit einem erheblichen Verlust an Flexibilität verbunden ist.

Beitrag B2 „Dynamic programming decomposition for choice-based revenue management with flexible products“ knüpft an die oben beschriebene Problematik an. Die Grundidee ist es, das zugrundeliegende DP so zu reformulieren, dass der Zustandsraum anstelle der Reservierungen von flexiblen Produkten die Restkapazität von neu geschaffenen, künstlichen Ressourcen umfasst. Im Beitrag wird ein entsprechendes Vorgehen entwickelt und die Äquivalenz der beiden DPs gezeigt. Das Vorgehen basiert auf der Anwendung der Fourier-Motzkin-Elimination auf das im DP inhärente Zulässigkeitsproblem (vgl. Schrijver (1998), Kap. 12.2 für eine Einführung in Fourier-Motzkin-Elimination). Zudem erfolgt für gängige Anwendungsfälle eine Analyse von Struktur sowie Größe der resultierenden, künstlichen Ressourcennetzwerke. Der vorgeschlagene Ansatz hat den Vorteil, dass Standardmethoden des Revenue Managements ohne Anpassung angewendet werden können. Dabei geht – insbesondere im Gegensatz zu Heuristiken mit sofortiger bzw. temporärer Zuweisung – der Vorteil der zeitlichen Flexibilität nicht verloren. In einer Simulationsstudie werden die damit einhergehenden Erlössteigerungen am Beispiel von Dekompositionsansätzen, die sich in die Klasse der Wertfunktionsapproximationen einordnen lassen, quantifiziert (vgl. Talluri und van Ryzin (2004b), Kap. 3.4, sowie Liu und van Ryzin (2008) für eine Beschreibung der im Revenue Management üblichen Dekompositionen von DPs).

Beitrag B3 “Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches“ quantifiziert zunächst explizit den Wert der zeitlichen Flexibilität. Anschließend wird gezeigt, dass in den bisher üblichen Erweiterungen gängiger statischer, linearer Ersatzmodelle dieser sog. Wert der Flexibilität unterschätzt wird. Die direkte Verwendung der Ersatzmodelle zur Kapazitätssteuerung führt damit tendenziell zu einem zu geringen Absatz von flexiblen Produkten. Im Beitrag wird dieses Problem durch eine systematische, künstliche Erhöhung des Verkaufserlöses von flexiblen Produkten „behooben“. Die Kalibrierung der künstlichen Erhöhung erfolgt mit Standardmethoden der simulationsbasierten Optimierung.

Beitrag B4 “An EMSR-based approach for revenue management with integrated upgrade decisions“ betrachtet ausschließlich Upgrades im Single-Leg-Fall. Dies erfordert

jedoch – im Gegensatz zum klassischen Revenue Management – die Betrachtung mehrerer, gemäß einer Upgrade-Hierarchie geordneter Ressourcentypen. In der Passagierluftfahrt bezeichnet man diese Ressourcentypen bspw. als Beförderungsklassen. Für den betrachteten Anwendungsfall zeigen Steinhardt und Gönsch (2012), dass die Integration der zeitlichen Flexibilität keine Erlössteigerungen ermöglicht. Stattdessen kann eine Reservierung immer unmittelbar im Verkaufszeitpunkt dem niedrigsten möglichen Ressourcentyp zugewiesen werden. Auf dieser Eigenschaft aufbauend wird einer Erweiterung des bekannten EMSR-a Verfahrens, das ursprünglich nur für den Fall einer Resource entwickelt wurde, zur Herleitung von Schutzlimits vorgestellt.

I.4 Revenue Management bei Risikoaversion

Wie in Kapitel I.1 angesprochen, optimieren Ansätze des Revenue Managements in der Regel den Erwartungswert des im Buchungshorizont erzielten Erlöses. Die Annahme eines risikoneutralen Entscheiders ist dabei in traditionellen Anwendungsfeldern meist gerechtfertigt. Aufgrund der großen Anzahl ähnlicher Entscheidungsprobleme beeinflusst ein einzelnes Verkaufsereignis bzw. ein einzelner Buchungshorizont den Unternehmenserfolg nur unwesentlich und das Gesetz der großen Zahlen garantiert die Konvergenz des durchschnittlichen Erlöses hin zu dessen Erwartungswert.

Eine risikoaverse Bewertung hingegen liegt nahe, wenn Buchungsprozesse nur selten stattfinden und/oder das Ergebnis eines Buchungsprozesses erfolgskritisch für das Unternehmen ist (vgl. Barz (2007)). Beispielhaft genannt wird meist ein Konzertveranstalter, der nur wenige große Konzerte pro Jahr organisiert (vgl. Levin et al. (2008)). Zudem haben Studien gezeigt, dass menschliche Entscheidungsträger insbesondere bei kleineren Unternehmen die risikoneutralen Empfehlungen von Revenue Management-Systemen als zu riskant wahrnehmen und diese daher häufig verwerfen (vgl. Barz (2007) oder Singh (2011)).

Bisher existieren im Revenue Management vergleichsweise wenige risikoaverse Ansätze. Die Berücksichtigung von Risikoaversion im Rahmen der dynamischen Optimierung erfordert dabei häufig eine Erweiterung des Zustandsraums um den im Buchungsverlauf

kumulierten Erlös oder ähnliche Erweiterungen, so dass das resultierende DP nicht mehr mit Standardmethoden (vgl. Kapitel I.1) gelöst werden kann. Zudem ist die aus risikoaversen DPs resultierende Kapazitätssteuerung in der Regel nicht kompatibel mit bestehenden Revenue Management-Systemen.

Beitrag B5 „Practical decision rules for risk-averse revenue management using simulation-based optimization“ versucht, diese Forschungslücke zu schließen. Der vorgeschlagene Ansatz beruht auf kleinen, intuitiven Modifikationen ursprünglich risikoneutraler Methoden der Kapazitätssteuerung über kalibrierbare Parameter und ist deshalb einfach in bestehende Revenue Management-Systeme zu integrieren. Die Parameter können mit Standardmethoden der simulationsbasierten Optimierung gelernt werden, so dass der Ansatz weder ein bestimmtes Risikomaß noch ein bestimmten Nachfragemodell voraussetzt. Die Auswirkungen des Ansatzes auf die Ergebnisse einer Kapazitätssteuerung werden am Beispiel der Erwartungsnutzentheorie sowie des zur Abschätzung von Finanzrisiken bekannten Conditional Value-at-Risk veranschaulicht.

Literaturverzeichnis

- Adelman D (2007) Dynamic bid prices in revenue management. *Operations Research* 55(4): 647–661.
- Agatz N, Campell AM, Fleischmann M, van Nunen J, Savelsbergh M (2013) Revenue management opportunities for internet retailers. *Journal of Revenue and Pricing Management* 12(2): 128–138.
- Bartodziej P, Derigs U, Zils M (2007) O&D revenue management in cargo airlines – a mathematical programming approach. *OR Spectrum* 29(1): 105–121.
- Barz C (2007) Risk-averse capacity control in revenue management. Springer, Berlin.
- Belobaba PP (1987) Air travel demand and airline seat inventory management. Dissertation, Massachusetts Institute of Technology, Cambridge.
- Belobaba PP (1989) Application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37(2): 183–197.
- Belobaba PP (1992) Optimal vs. heuristic methods for nested seat allocation. Vortrag bei der AGIFORS Reservations and Yield Management Study Group, Brüssel.
- Brooks C (2014) Introductory econometrics for finance. 3. Aufl. Cambridge University Press, Cambridge.
- Domschke W, Drexl A, Klein R, Scholl A (2015) Einführung in Operations Research. 9. Aufl. Springer, Berlin.
- Fink A, Reinert T (2006) Modeling and solving the short-term car rental logistics problem. *Transportation Research Part E: Logistics and Transportation Review* 42(4): 272–292.
- Gallego G, Iyengar G, Phillips RL, Dubey A (2004) Managing flexible products on a network. Arbeitspapier, Columbia University, New York.
- Gallego G, Phillips RL (2004) Revenue management of flexible products. *Manufacturing & Service Operations Management* 6(4): 321–337.
- Gönsch J, Steinhardt C (2015) On the incorporation of upgrades into airline network revenue management. *Review of Managerial Science* 9(4): 635–660.
- Gühlich H, Fleischmann M, Mönch L, Stolletz R (2015) Revenue management in a multi-stage make-to-order production system. Arbeitspapier, Universität Mannheim, Mannheim.
- Hillier FS, Lieberman GJ (2010) Introduction to operations research. 9. Aufl. Mc-Graw-Hill, New York.
- Jerath K, Netessine S, Veeraraghavan SK (2010) Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Science* 56(3): 430–448.

- Kimms A, Müller-Bungart M (2007) Revenue management for broadcasting commercials: The channel's problem of selecting and scheduling advertisements to be aired. *International Journal of Revenue Management* 1(1): 28–44.
- Klein R (2001) Revenue Management: Quantitative Methoden zur Erlösmaximierung in der Dienstleistungsproduktion. *Betriebswirtschaftliche Forschung und Praxis* 53(3): 245–259.
- Klein R, Steinhardt C (2008) Revenue Management – Grundlagen und Mathematische Methoden. Springer, Berlin.
- Kök AG, Fisher ML, Vaidyanathan R (2009) Assortment planning: Review of literature and industry practice. In: Agrawal N, Smith SA (eds) Retail supply chain management: Quantitative models and empirical studies. Springer, New York: 99–153.
- Levin Y, McGill JI, Nediak M (2008) Risk in revenue management and dynamic pricing. *Operations Research* 56(2): 326–343.
- Liu Q, van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2): 288–310.
- McFadden D (2001) Economic choices. *American Economic Review* 91(3): 351–378.
- Meissner J, Strauss A (2012) Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research* 216(2): 459–468.
- Miranda Bront JJ, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Operations Research* 57(3): 769–784.
- Müller-Bungart M (2007) Revenue management with flexible products: Models and methods for the broadcasting industry. Springer, Berlin.
- Nickel S, Stein O, Waldmann K (2011) *Operations Research*. Springer, Berlin.
- Petrick A, Gönsch J, Steinhardt C, Klein R (2010) Dynamic control mechanisms for revenue management with flexible products. *Computers & Operations Research* 37(11): 2027–2039.
- Petrick A, Steinhardt C, Gönsch J, Klein R (2012) Using flexible products to cope with demand uncertainty in revenue management. *OR Spectrum* 34(1): 215–242.
- Phillips RL (2005) *Pricing and revenue optimization*. Stanford University Press, Stanford.
- Powell WB (2011) *Approximate dynamic programming: Solving the curses of dimensionality*. 2. Aufl. Wiley, Hoboken, New Jersey.
- Prokopyev OA, Huang H, Pardalos PM (2005) On complexity of unconstrained hyperbolic 0–1 programming problems. *Operations Research Letters* 33(3): 312–318.

- Puterman ML (2005) Markov decision processes: Discrete stochastic dynamic programming. Wiley, Hoboken, New Jersey.
- Schrijver A (1998) Theory of linear and integer programming. Wiley, Chichester.
- Singh P (2011) Risk in revenue management. Acknowledging risk's existence and knowing how to minimize it. *Inform Analytics Magazine* May/June: 16–19.
- Steinhardt C, Gönsch J (2012) Integrated revenue management approaches for capacity control with planned upgrades. *European Journal of Operational Research* 223(2): 380–391.
- Subramanian J, Stidham Jr. S, Lautenbacher CJ (1999) Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science* 33(2): 147–167.
- Talluri KT, van Ryzin G (2004a) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.
- Talluri KT, van Ryzin G (2004b) The theory and practice of revenue management. Springer, New York.
- Train K (2009) Discrete choice methods with simulation. 2. Aufl. Cambridge University Press, Cambridge.
- Yang X, Strauss AK, Currie CSM, Eglese R (2016) Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science* 50(2), 473–488.

II Beiträge

Beitrag B1:

Least squares approximate policy iteration for learning bid prices in choice-based revenue management

Sebastian Koch (Universität Augsburg)

Computers & Operations Research 77(2017): 240–253.

DOI: <https://doi.org/10.1016/j.cor.2016.07.015>

Beitrag B2:

Dynamic programming decomposition for choice-based revenue management with flexible products

Sebastian Koch ^a, Jochen Gönsch ^b, Claudius Steinhardt ^c

(^a Universität Augsburg, ^b Universität Duisburg-Essen, ^c Universität der Bundeswehr München)

Erscheint in: Transportation Science.

DOI: <https://doi.org/10.1287/trsc.2017.0743>

Abstract

We reconsider the stochastic dynamic program of revenue management with flexible products and customer choice behavior as proposed by Gallego et al. [Gallego G, Iyengar G, Phillips RL, Dubey A (2004) Managing flexible products on a network. Working paper, Columbia University, New York]. In the scientific literature on revenue management, as well as in practice, the prevailing strategy to operationalize dynamic programs is to decompose the network by resources and solve the resulting one-dimensional problems. However, to date, these dynamic programming decomposition approaches have not been applicable to problems with flexible products, because sold flexible products must be included in the dynamic program's state space and do not correspond directly to resources.

In this paper, we contribute to the existing research by presenting a general approach to operationalizing revenue management with flexible products and customer choice in a dynamic programming environment. In particular, we reformulate the original dynamic program by means of Fourier-Motzkin elimination to obtain an equivalent dynamic program with a standard resource-based state space. This reformulation allows the application of dynamic programming decomposition approaches. Numerical experiments show that the new approach has a superior revenue performance and that its average revenues are close to the upper bound on the optimal expected revenue from the choice-based deterministic linear program (CDLP). Moreover, our reformulation improves the revenues by up to 8% compared to an extended variant of a standard choice-based approach without flexible products that immediately assigns flexible products after their sale.

Keywords: Revenue Management, Flexible Products, Dynamic Programming Decomposition, Customer Choice, Fourier-Motzkin Elimination

1 Introduction

While the specification of common, regular products is fixed in advance, a flexible product consists of two or more alternative specifications, such that the seller will assign the purchaser to one of these alternatives at a later point in time (see, e.g., Gallego et al. (2004)). From a revenue management point of view, this supply-side flexibility leads to improved capacity utilization and mitigates the negative impact of forecast errors that often occur due to demand's stochastic nature. From a marketing perspective, flexible products are an interesting tool for market segmentation. Owing to their inherent uncertainty, and because they are offered at a lower price, they are perceived as inferior by the customer and induce additional low value demand, while avoiding excess cannibalization. Flexible products have to be distinguished from opaque products, whose utilized resources the firm determines immediately after the sale, thus losing the benefit of postponing the products' assignment.

In this paper, we reconsider the problem of choice-based revenue management with flexible products, which Gallego et al. (2004) introduced. In their paper, these authors incorporated flexible products into the dynamic program (DP) for revenue management with arbitrary resource networks, while assuming choice-based demand behavior. To incorporate flexible products, they extended the state space of the DP from a purely resource-based one to a space that contains resources' remaining capacity as well as commitments reflecting sold flexible products that must be assigned to alternatives later on. An inherent feasibility problem ensures that the remaining capacity can satisfy the commitments.

In the traditional setting without flexible products, the standard way to make such a multidimensional DP operational is through dynamic programming decomposition (DPD). Standard DPD can be roughly summarized as follows: At first, a linear approximation of the corresponding DP is solved to obtain dual variables which capture the network effects. The dual variables are then used to decompose the network by resources. Doing so provides single-resource DPs that are easily solved to optimality. Within the single-resource DPs, only products that use the corresponding resource are

considered and the capacity consumption of other resources is captured by reducing the products' revenues according to the dual variables from the linear approximation.

However, this approach is not applicable to the DP with flexible products, because its state space contains commitments that do not correspond directly to resources. In other words, if a flexible product is sold, the resources whose capacity is consumed are unknown, as they are determined later. Thus, an assignment to a single-resource DP and the immediate reduction of the remaining capacity are not possible if the flexibility should be preserved.

Our main contribution is that we show how to obtain an equivalent reformulation of the original DP whose state space is no longer based on commitments. The central idea is to apply Fourier-Motzkin elimination (FME; see, e.g., Schrijver (1998), Chapter 12.2) to the feasibility problem inherent in the DP. In doing so, additional “artificial” resources are added, allowing the flexible products to directly correspond to the artificial resources. This allows the reformulation as a standard revenue management problem without flexible products. We call this the surrogate approach. The key benefit of the new state space is that it enables the application of DPD and other standard methods, which make dynamic programming-based, large-scale implementations operational.

The remainder of this paper is structured as follows: In Section 2, we review the relevant scientific literature and position our work. In Section 3, we briefly summarize the standard DP of revenue management with flexible products, which Gallego et al. (2004) proposed, and restate the relevant notation. On this basis, we derive the surrogate reformulation in Section 4. By using numerical experiments, we evaluate DPD of the surrogate reformulation in Section 5. We use the upper bound from the optimal objective value of the choice-based deterministic linear program (CDLP), as well as two standard methods—the CDLP's primal solution and an adequate extension of the DPD ad hoc approach known from the literature—as benchmarks. In Section 6, we discuss our results and conclude the paper.

2 Related Literature

Initially, revenue management research assumed that demand was independent of the available products and of other customers (the well-known independent demand assumption). A considerable amount of work was done on the solution of single-resource problems (see, e.g., Littlewood (1972) for the earliest work; Belobaba (1987, 1989) for the expected marginal seat revenue heuristic; Lee and Hersh (1993), as well as Lautenbacher and Stidham (1999), for analyses of the exact DP formulation). However, as soon as networks of resources are considered, the corresponding DP formulations are difficult to solve even for small instances. Consequently, many heuristic approaches have been developed to approximate the DP. These approaches are mainly based on the idea of decomposing the network problem into a collection of smaller sub-problems. Decomposition is usually done by resources, while network effects are considered by adequately modifying the revenues of products that use more than one resource. The idea is to subsequently apply single-resource methods to the obtained sub-problems. We refer to Talluri and van Ryzin (2004b), Chapter 3.4, for an overview. Well-known decomposition approaches are origin-destination factor methods, fare proration (see, e.g., Kemmer et al. (2011) for an extension of the standard approach to large-scale applications), and DPD, which is most common in practice and literature, and is this paper's focus. The idea of DPD is to use the dual variables of a corresponding deterministic linear program (DLP) (see, e.g., Talluri and van Ryzin (1998)) to capture network effects and modify the products' revenues in the sub-problems. A further refinement of the standard DPD approach is studied in Zhang (2011). In addition, there are a few specific decomposition ideas (see, e.g., Cooper and Homem-de-Mello (2007), who include ideas from mathematical programming, and Birbil et al. (2014), who decompose the network by product types that require the same combination of resources).

Over the past decade, two major trends have emerged, which we will address in the following: first, the incorporation of demand-side substitution, which stems from customer choice behavior; second, the integration of supply-side substitution via flexible products.

Regarding the first trend, Talluri and van Ryzin (2004a) and Gallego et al. (2004) overcome the assumption of independent demand by considering customer choice behavior in the context of a single resource and a network of resources, respectively. In order to have a counterpart to the traditional DLP, Gallego et al. (2004) formulated the now well-known CDLP as a linear approximation of the underlying DP. Liu and van Ryzin (2008) and Miranda Bront et al. (2009) analyze the CDLP further. They assume that customer segments follow a standard multinomial logit model (see, e.g., Train (2009), Chapter 3) and that these segments consider buying products from disjoint and overlapping consideration sets, respectively. Gallego et al. (2015) reformulate the CDLP in respect of disjoint consideration sets; this reformulation avoids the exponential number of variables. Meissner et al. (2013) and Strauss and Talluri (2015) investigate weaker, but more efficient, deterministic linear approximations than the CDLP. Recent research has also examined many different customer choice models (see, e.g., Davis et al. (2014) for the nested logit model; Hosseinalifam (2014), Chapter 3, for a ranking-based customer choice model). Analogously to the independent demand setting, the CDLP is then used within an appropriate DPD approach. Liu and van Ryzin (2008) were the first to adapt the standard DPD to the choice-based setting. A number of subsequent papers have investigated this approach further (see Miranda Bront et al. (2009), as well as Kunnumkal and Topaloglu (2010), for a refinement of the standard DPD and derivations of upper bounds, respectively; Zhang and Adelman (2009), as well as Vossen and Zhang (2015b), for derivations of upper bounds and connections of DPD and the linear programming approach for approximate dynamic programming).

The second trend, that is, the consideration of flexible products, is also rooted in Gallego et al. (2004). These authors present a generalized DP formulation for flexible products in arbitrary resource networks that extends the state space by commitments. Their formulation, which also incorporates customer choice behavior, is standard for revenue management with flexible products today (see Section 3). However, subsequent research largely continued to follow the independent demand assumption. Among others, flexible products were investigated in the context of passenger aviation (see, e.g.,

Gallego and Phillips (2004)), air cargo revenue management (see, e.g., Bartodziej et al. (2006)), and the broadcasting industry (see, e.g., Kimms and Müller-Bungart (2007)). Upgrades can be seen as a special case of flexible products with hierarchically ordered alternatives (see, e.g., Gallego and Stefanescu (2009)).

Overbooking problems with no shows are somehow also related, because there are also commitments in the state space of the corresponding DP formulations. At the end of the booking horizon, an optimization problem is solved to determine which reservations should be denied, which is similar to the feasibility problem inherent in the DP formulation with flexible products. However, in overbooking, commitments correspond directly to resources, which is similar to the traditional revenue management setting without flexible products. Erdelyi and Topaloglu (2010) are thus able to separate the optimization problem at the end of the booking horizon by resources, thus making standard DPD applicable. Similarly, Erdelyi and Topaloglu (2009) and Kunnumkal and Topaloglu (2008) approximate the optimization problem at the end of the booking horizon with a function that is separable by products reflecting the individual overbooking costs. In doing so, the authors are able to decompose the DP by products.

In contrast, in revenue management with flexible products, the standard decomposition by resources as in DPD is not possible, because the products do not correspond directly to resources. Basically, two literature streams tackling this issue have emerged:

- In the first stream, the supply-side flexibility is relinquished, allowing a flexible product to actually become an opaque product. Technically, the flexible product is irrevocably assigned immediately after the sale to one of the alternatives. We refer to Talluri (2001) and Chen et al. (2010) for revenue management with opaque products. In doing so, the need to store a commitment for later assignment is eliminated, and the solely resource-based state space is retained, which renders DPD possible again (see Gönsch and Steinhardt (2013) for DPD with opaque products). Other authors have sought to at least partially retain the flexibility. For example, Petrick et al. (2010, 2012) use bid prices from a deterministic linear programming (DLP) for-

mulation and reassigned the sold flexible products when the DLP is resolved during the booking horizon.

- In the second stream, the supply-side flexibility is preserved at the cost of a restriction to special network structures. Often, only parallel resources and just one flexible product are considered (see, e.g., Gallego and Phillips (2004) and Oosten (2004)). Gönsch and Steinhardt (2015) also considered DPD approaches, but restrict themselves to independent demand and airline upgrading, where hierarchical upgrades can be granted independently on each leg of a multi-leg flight. They use results from production planning that Leachman and Carmon (1992) obtained earlier.

This paper overcomes the drawbacks inherent in both literature streams. It enables DPD under customer choice behavior with arbitrary network structures, while fully retaining the supply-side flexibility. Additionally, in order to obtain a valid benchmark procedure for comparison, we adequately adapt an existing approach from choice-based revenue management to the flexible products setting. Our benchmark approach follows the idea of the first literature stream described above. In particular, it incorporates flexible products into the DPD approach of Liu and van Ryzin (2008) by immediately assigning them after sale.

3 Standard model formulation with flexible products

In the following, we first summarize the choice-based revenue management problem with flexible products (see Gallego et al. (2004)) and repeat the relevant notation (Section 3.1). Thereafter, we restate the corresponding DP (Section 3.2).

3.1 Problem formulation and notation

We consider a firm that sells regular products $j \in \mathcal{J} = \{1, \dots, n^{reg}\}$ and flexible products $k \in \mathcal{K} = \{1, \dots, n^{flex}\}$. These products use resources $h \in \mathcal{H} = \{1, \dots, m\}$ jointly and may be linked to sale restrictions or rules in order to segment the market. The customers arrive successively and stochastically over time before service provision. The

regular products are associated with revenues $\mathbf{r}^{reg} = (r_1^{reg}, \dots, r_n^{reg})^T$. Furthermore, each regular product j has a capacity consumption $\mathbf{a}_j = (a_{1j}, \dots, a_{mj})^T$, which is $a_{hj} = 1$ if product j uses resource h , and $a_{hj} = 0$ otherwise. Regarding the flexible products with revenues $\mathbf{r}^{flex} = (r_1^{flex}, \dots, r_n^{flex})^T$, the resources to be utilized can be decided just before service provision. More precisely, a customer who buys flexible product k is guaranteed the resources \mathbf{a}_j of one of the alternative regular products $j \in \mathcal{M}_k \subseteq \mathcal{J}$. The obtained revenue r_k^{flex} is fixed in advance and independent of this assignment.

For notational convenience, we use $(\mathbf{A}, \mathcal{M})$ to denote the network structure, where $\mathbf{A} = [a_{hj}]_{m \times n^{reg}}$ is the regular products' capacity consumption matrix and, by slightly abusing notation, $\mathcal{M} = (\mathcal{M}_1, \dots, \mathcal{M}_{n^{flex}})^T$ is the flexible products' vector of alternatives.

The state of the selling process is described by the remaining capacity $\mathbf{c} = (c_1, \dots, c_m)^T$ and the vector of commitments $\mathbf{y} = (y_1, \dots, y_{n^{flex}})^T$, which denotes the number of sold flexible products. Selling a regular product j reduces the remaining capacity to $\mathbf{c} - \mathbf{a}_j$, and selling a flexible product k increases the commitment vector to $\mathbf{y} + \mathbf{e}_k$, with \mathbf{e}_k referring to the k -th standard basis vector in $\mathbb{R}^{n^{flex}}$.

We discretize the booking horizon into T time periods, such that in each period t there is, at most, one customer arrival. The periods are numbered backward in time, and w.l.o.g., the probability λ of a customer's arrival is time-homogeneous. Any capacity remaining at the end of the booking horizon is worthless and overbooking of the given resources' capacity is not allowed. In each period t , the firm's decision problem is to determine a subset of products to offer, called the offer set. Given an offer set $S \subseteq \mathcal{J} \cup \mathcal{K}$, an arriving customer purchases product j with probability $P_j^{reg}(S)$, product k with probability $P_k^{flex}(S)$, and makes no purchase with probability $P_0(S)$. The firm aims to maximize its total overall revenue.

In what follows, we omit the index sets of the products and resources where possible. For example, the notation $\forall k$ means $\forall k \in \mathcal{K}$, \sum_k means $\sum_{k \in \mathcal{K}}$, and \max_S means $\max_{S \subseteq \mathcal{J} \cup \mathcal{K}}$.

3.2 Dynamic programming formulation

Given the remaining capacity \mathbf{c} and the commitments \mathbf{y} , the optimal expected revenue-to-go with t time periods left is denoted by $V_t(\mathbf{c}, \mathbf{y})$ and satisfies the Bellman equation (**DP-flex**)

$$\begin{aligned} V_t(\mathbf{c}, \mathbf{y}) = \max_S \{ & \sum_j \lambda \cdot P_j^{reg}(S) \cdot (r_j^{reg} + V_{t-1}(\mathbf{c} - \mathbf{a}_j, \mathbf{y})) \\ & + \sum_k \lambda \cdot P_k^{flex}(S) \cdot (r_k^{flex} + V_{t-1}(\mathbf{c}, \mathbf{y} + \mathbf{e}_k)) \\ & + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}(\mathbf{c}, \mathbf{y}) \} \end{aligned} \quad (1)$$

with the boundary conditions $V_t(\mathbf{c}, \mathbf{y}) = -\infty$ if $(\mathbf{c}, \mathbf{y}) \notin \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ and $V_0(\mathbf{c}, \mathbf{y}) = 0$ if $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$.

The condition $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ describes a feasible state and holds if the capacity is non-negative and can satisfy all commitments in the given network structure $(\mathbf{A}, \mathcal{M})$. More formally, $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ if and only if there exist (nonnegative and integer) distribution variables y_{kj} denoting how many commitments regarding flexible product k will be fulfilled with alternative j satisfying the **feasibility problem** (see Gallego et al. (2004)):

$$\sum_k \sum_{j \in \mathcal{M}_k} a_{hj} \cdot y_{kj} \leq c_h \quad \forall h \quad (2)$$

$$\sum_{j \in \mathcal{M}_k} y_{kj} = y_k \quad \forall k \quad (3)$$

$$y_{kj} \in \mathbb{Z}^+ \quad \forall k, j \in \mathcal{M}_k \quad (4)$$

To illustrate the problem, we introduce the following running example (expressed in airline terminology) that will also be reconsidered in Section 4 to illustrate the reformulation as well as the transformation we propose.




Figure 1: Airline network in running example

Example: Consider an airline that offers transportation from A to B over a hub H at two different times of day as depicted in Figure 1, resulting in a small network with $m = 4$ legs. There is one flexible product $k = 1$ that uses either legs $h = 1$ and $h = 2$ (alternative y_{11}) or legs $h = 3$ and $h = 4$ (alternative y_{22}). Table 1 represents the corresponding instance of the feasibility problem (constraints (2) and (3)) without integrality constraints, where the first column refers to the row number (rows (i)–(iv) are instances of constraint (2), row (v) is an instance of constraint (3)):

Row	y_{11}	y_{12}		Right-hand side
(i)	1		\leq	c_1
(ii)	1		\leq	c_2
(iii)		1	\leq	c_3
(iv)		1	\leq	c_4
(v)	1	1	$=$	y_1

Table 1: Feasibility problem in running example

4 Surrogate approach

The commitments in the state space of DP-flex (1) are obviously necessary to solve the feasibility problem (2)–(4) throughout the booking horizon. However, because the commitments do not correspond directly to resources, they inhibit the use of decomposition by resources.

To overcome this problem, we suggest applying FME in order to project the distribution variables y_{kj} out of the feasibility problem. In doing so, additional “artificial” resources are added, with the flexible products now corresponding directly to the artificial resources (Section 4.1). In Section 4.2, we use the DP formulation to show that this allows the reformulation as a standard revenue management problem without flexible products,

such that standard solution approaches and heuristics can be used. In Section 4.3, we analyze the problem size of several network types in which flexible products occur in practice.

4.1 Transformation of the feasibility problem

In this subsection, we focus on the feasibility problem (2)–(4) by thinking of it as a static problem that must be solved for a given network $(\mathbf{A}, \mathcal{M})$ at some point in time during the booking horizon, in order to decide whether a state (\mathbf{c}, \mathbf{y}) is feasible. We show how the distribution variables can be eliminated and explain the output of this elimination process.

In the feasibility problem, the integrality of the assignments of customers to alternatives is ensured by $y_{kj} \in \mathbb{Z}^+$ in (4). However, when projecting out the distribution variables by means of FME, we cannot keep this constraint and need a formulation that includes only \leq constraints:

$$\sum_k \sum_{j \in \mathcal{M}_k} a_{hj} \cdot y_{kj} \leq c_h \quad \forall h \quad (5)$$

$$\sum_{j \in \mathcal{M}_k} -y_{kj} \leq -y_k \quad \forall k \quad (6)$$

$$-y_{kj} \leq 0 \quad \forall k, j \in \mathcal{M}_k \quad (7)$$

Constraints (5)–(7) are the linear relaxation of (2)–(4). Whereas (5) and (7) correspond directly to (2) and (the relaxed) (4), constraints (6) may be less obvious. They follow from rewriting (3) as $\sum_{j \in \mathcal{M}_k} y_{kj} \geq y_k \forall k$, which is equivalent, because the feasible region defined by (2) and (the relaxed) (4) with regard to y_{kj} is a convex polytope including $\mathbf{0}$. In order to ensure that (2)–(4) can technically be replaced with (5)–(7), we claim that the following condition needs to hold:

Condition 1: If the linear relaxation (5)–(7) has an arbitrary solution, there exists also an integer solution (i.e., a solution given the same number of commitments which (additionally) satisfies (4)).

Please note that Condition 1 is satisfied in most applications. For example, it is a sufficient condition for Condition 1 to be valid that the left-hand side coefficient matrix of the feasibility problem (2)–(3) (or, equivalently, (5)–(6)) is totally unimodular, since adding an identity matrix as in (7) would preserve this property (see, e.g., Martin (1999), Chapter 14.2). Total unimodularity is, at least for relevant problem sizes, easy to check (see, e.g., Walter and Truemper (2013)). Please note that total unimodularity in our case refers to the complete feasibility problem including (3) and thus to a different matrix than in the common discussion in the revenue management literature, where it refers to the left-hand side matrix of the well-known deterministic linear program (DLP; see, e.g., Talluri and van Ryzin (2004b)), that is, (2) without (3) but with an additional identity matrix resulting from demand constraints. It is well-known that total unimodularity is satisfied, for example, in problems for which corresponding network flow formulations can be constructed. For our feasibility problem, the construction of such network flow formulations can be performed analogously to the construction for the DLP (see, e.g., Glover et al. (1982), as well as Bertsimas and Popescu (2003) for the construction of network flow formulations in origin-destination networks, and Chen (1998) for hotel networks). Even more, recall that total unimodularity is only a sufficient condition, and there are also many other settings without total unimodularity which satisfy Condition 1.

Note that in the case that Condition 1 does not hold, network instances whose capacity is slightly overestimated could potentially result from the linear relaxation (see also Proposition 1 and the note thereafter).

Now, we can project the distribution variables y_{kj} out of (5)–(7) by using FME. The classical FME idea can be summarized as follows: Consider that we want to project variable x out of the inequality system $LB_i \leq a_i \cdot x \forall i, b_j \cdot x \leq UB_j \forall j$. This inequality system has a feasible solution if and only if $\max_i \frac{LB_i}{a_i} \leq \min_j \frac{UB_j}{b_j}$, which is equal to the system of linear inequalities $b_j \cdot LB_i \leq a_i \cdot UB_j \forall i, j$. Therefore, the initial inequality system can be replaced equivalently by the latter constraints.

We can now apply this idea to our setting, considering one distribution variable after another. The important point here is that we also treat the state of the selling process as variables. To formalize this approach, let **LHS** and **RHS** be the left-hand side and right-hand side coefficient matrices of (5)–(7), respectively. With the $\sum_k |\mathcal{M}_k| \times 1$ vector of distribution variables denoted by $\mathbf{y}^{kj} = (y_{kj})_{\forall k, j \in \mathcal{M}_k}$, (5)–(7) can be rewritten as:

$$\mathbf{LHS} \cdot \mathbf{y}^{kj} \leq \mathbf{RHS} \cdot (\mathbf{c}^T | \mathbf{y}^T | \mathbf{0}^T)^T \quad (8)$$

Please note that **LHS** and **RHS** only depend on the network structure $(\mathbf{A}, \mathcal{M})$ and are the coefficients. Now, Algorithm 1 creates a projection of (8) by applying a sequence of FMEs to the distribution variables.

Algorithm 1: Elimination of distribution variables from the feasibility problem

```

1  for  $d = 1$  to  $\sum_k |\mathcal{M}_k|$  do                                ▷ consider one distribution variable  $y_{kj}$ 
2       $nRows \leftarrow$  number of rows of LHS
3       $Pos \leftarrow \{row \in \{1, \dots, nRows\}: lhs_{row,d} > 0\}$           ▷ partition row indices
4       $Neg \leftarrow \{row \in \{1, \dots, nRows\}: lhs_{row,d} < 0\}$ 
5       $Null \leftarrow \{row \in \{1, \dots, nRows\}: lhs_{row,d} = 0\}$ 
6       $nRows \leftarrow |Null \cup (Pos \times Neg)|$ 
                                   ▷ number of rows after eliminating the current distribution variable
7  Let  $biject$  be a bijection that maps  $\{1, \dots, nRows\}$  onto  $Null \cup (Pos \times Neg)$ 
                                   ▷  $biject$  is an arbitrary indexation of the new rows after the elimination
8  for  $row = 1$  to  $nRows$  do                                    ▷ construct a new row
9      if  $biject(row) \in Null$  then                             ▷ copy row without change
10          $\mathbf{LHS}_{row}^{new} \leftarrow \mathbf{LHS}_{biject(row)}$  and  $\mathbf{RHS}_{row}^{new} \leftarrow \mathbf{RHS}_{biject(row)}$ 
11     else                                                    ▷  $biject(row) \in Pos \times Neg$  and add these rows
12          $(pos, neg) \leftarrow biject(row)$ 
13          $\mathbf{LHS}_{row}^{new} \leftarrow lhs_{neg,d} \cdot \mathbf{LHS}_{pos} + lhs_{pos,d} \cdot \mathbf{LHS}_{neg}$ 
14          $\mathbf{RHS}_{row}^{new} \leftarrow lhs_{neg,d} \cdot \mathbf{RHS}_{pos} + lhs_{pos,d} \cdot \mathbf{RHS}_{neg}$ 
15     Set  $\mathbf{LHS} \leftarrow \mathbf{LHS}^{new}$  and  $\mathbf{RHS} \leftarrow \mathbf{RHS}^{new}$ 
16 return RHS                                                ▷ return only RHS, because  $\mathbf{LHS} = \mathbf{0}$ 

```

It can briefly be explained as follows: The distribution variables are projected out step-wise (line 1), considering one column (i.e., distribution variable) of **LHS** after another. In an iteration, the row indices of the current feasibility problem are partitioned by their coefficient of **LHS** into sets of rows with positive, negative, and null coefficients (lines 2–5). Based on this, the new number of rows (after eliminating the current distribution variable) is determined (line 6), and an arbitrary indexation of these rows is introduced

(line 7). Finally, in lines 8–15, the current distribution variable is projected out. Please note that \mathbf{LHS}_{row} (\mathbf{RHS}_{row}) refers to the row -th row of \mathbf{LHS} (\mathbf{RHS}).

Additionally, as pointed out by Bertsimas and Tsitsiklis (1997), Chapter 2.8, redundant rows should be regularly eliminated while performing FME (see, e.g., Paulraj and Sumathi (2010) on finding redundant constraints in linear inequality systems).

Example (cont'd): Returning to the example from Section 3.2, the algorithm is illustrated by Table 2. The first set of rows (i)–(vii) refers to the initial feasibility problem, the second (i')–(vii') to the feasibility problem after projecting out y_{11} , and the third (i'')–(viii'') to the transformed feasibility problem after projecting out y_{12} . The last column refers to the operation performed to obtain the row.

Row	y_{11}	y_{12}		Right-hand side	Operation
(i)	1		\leq	c_1	
(ii)	1		\leq	c_2	
(iii)		1	\leq	c_3	
(iv)		1	\leq	c_4	
(v)	-1	-1	\leq	$-y_1$	
(vi)	-1		\leq	0	
(vii)		-1	\leq	0	
(i')		1	\leq	c_3	(iii)
(ii')		1	\leq	c_4	(iv)
(iii')		-1	\leq	0	(vii)
(iv')		-1	\leq	$c_1 - y_1$	(i) + (v)
(v')			\leq	c_1	(i) + (vi)
(vi')		-1	\leq	$c_2 - y_1$	(ii) + (v)
(vii')			\leq	c_2	(ii) + (vi)
(i'')			\leq	c_1	(v')
(ii'')			\leq	c_2	(vii')
(iii'')			\leq	c_3	(i') + (iii')
(iv'')			\leq	c_4	(ii') + (iii')
(v'')			\leq	$c_1 + c_3 - y_1$	(i') + (iv')
(vi'')			\leq	$c_1 + c_4 - y_1$	(ii') + (iv')
(vii'')			\leq	$c_2 + c_3 - y_1$	(i') + (vi')
(viii'')			\leq	$c_2 + c_4 - y_1$	(ii') + (vi')

Table 2: Algorithm 1 in running example

In the first set of rows, there are only \leq -constraints according to (5)–(7). Now, to project out y_{11} , the rows with null coefficient ($Null = \{(iii),(iv),(vii)\}$) are copied without

change (resulting in rows (i')–(iii')). The rows with positive coefficient ($Pos = \{(i),(ii)\}$) and negative coefficient ($Neg = \{(v),(vi)\}$) are added according to lines 13–14 of the algorithm (resulting in rows (iv')–(vii')). The second iteration is performed analogously, finally leading to the transformed feasibility problem given by rows (i'')–(viii''). The transformed feasibility problem obviously consists of non-negativity constraints of regular resources' remaining capacity, i.e., $0 \leq c_h \forall h = 1, \dots, 4$, as well as four additional constraints that each time check non-negativity of two regular resources' remaining capacity less the commitments for the flexible product:

$$\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \leq \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} \cdot c_1 + \begin{pmatrix} 0 \\ 0 \\ 1 \\ 1 \end{pmatrix} \cdot c_2 + \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix} \cdot c_3 + \begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} \cdot c_4 - \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \cdot y_1$$

In general, we obtain the following transformed feasibility problem (9) and (10) which has the same form as the feasibility check in the traditional setting without flexible products, in which only the non-negativity of the remaining capacity is checked (see, e.g., Talluri and van Ryzin (2004b), Chapter 3.2):

$$0 \leq c_h \quad \forall h \in \mathcal{H} \quad (9)$$

$$0 \leq \sum_h \tilde{f}_{ih} \cdot c_h - \sum_k \tilde{b}_{ik} \cdot y_k \quad \forall i \in \tilde{\mathcal{H}} = \{1, \dots, \tilde{m}\} \quad (10)$$

Constraints (9), in which the non-negativity of the regular resources' remaining capacity c_h is checked, are in fact identical. Constraints (10) can be interpreted as analogous conditions that require the non-negativity of some additional, artificial resources $\tilde{\mathcal{H}} = \{1, \dots, \tilde{m}\}$. Please note that the number \tilde{m} of artificial resources as well as the values of \tilde{f}_{ih} and \tilde{b}_{ik} are determined by FME. In Section 4.3, we investigate the number \tilde{m} subject to different network types.

Each artificial resource i has a capacity of $\tilde{c}_i = \sum_h \tilde{f}_{ih} \cdot c_h - \sum_k \tilde{b}_{ik} \cdot y_k$ and can be considered a pool of several regular resources that captures their alternative usage: It pools the capacity of some resources (those with $\tilde{f}_{ih} = 1$) and is required by regular products needing these resources, as well as by one or more flexible products that use these resources alternatively (those with $\tilde{b}_{ik} = 1$).

To ease notation in the following, we group the coefficients \tilde{f}_{ih} and \tilde{b}_{ik} into $\tilde{\mathbf{f}}_h = \tilde{\mathbf{f}}_h(\mathbf{A}, \mathcal{M})$ and $\tilde{\mathbf{b}}_k = \tilde{\mathbf{b}}_k(\mathbf{A}, \mathcal{M})$. Furthermore, we group the artificial resources represented by the right hand side of (10) into $\tilde{\mathbf{c}} = \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) = \sum_h \tilde{\mathbf{f}}_h(\mathbf{A}, \mathcal{M}) \cdot c_h - \sum_k \tilde{\mathbf{b}}_k(\mathbf{A}, \mathcal{M}) \cdot y_k$. Finally, we define $\tilde{\mathbf{a}}_j = \tilde{\mathbf{a}}_j(\mathbf{A}, \mathcal{M}) = \sum_h \tilde{\mathbf{f}}_h(\mathbf{A}, \mathcal{M}) \cdot a_{hj}$, which can be interpreted as a regular product's capacity consumption of artificial resources. Given these definitions, (9) and (10) can be abbreviated to

$$\mathbf{c} \geq \mathbf{0} \tag{11}$$

$$\tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) \geq \mathbf{0}, \tag{12}$$

and we can state the following result:

Proposition 1: Given Condition 1 holds, $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ (that is, the feasibility problem (2)–(4) has a solution) if and only if $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y})) \geq \mathbf{0}$ (that is, (11)–(12) has a solution).

Proof: See Appendix A.

Note that if Condition 1 does not hold, the transformed feasibility problem in (11)–(12) becomes heuristic, and the artificial resources slightly overestimate capacity for flexible product k by, at most, $|\mathcal{M}_k|$ capacity units.

Example (cont'd): To illustrate the notation in our example, there are $\tilde{m} = 4$ artificial resources. Resource $h = 1$ is included in artificial resources $i = 1$ and $i = 2$, ($\tilde{\mathbf{f}}_1 = (1, 1, 0, 0)^T$) and so on, that is, $\tilde{\mathbf{f}}_2 = (0, 0, 1, 1)^T$, $\tilde{\mathbf{f}}_3 = (1, 0, 1, 0)^T$, $\tilde{\mathbf{f}}_4 = (0, 1, 0, 1)^T$. The flexible product consumes capacity on all artificial resources ($\tilde{\mathbf{b}}_1 = (1, 1, 1, 1)^T$).

4.2 Reformulation as a standard revenue management problem

We next consider the dynamic revenue management problem again and show how Proposition 1 allows for managing the sales process of flexible products. A straightforward application of the transformed feasibility problem would be to replace (2)–(4) repeatedly to check which products can be offered for sale. More precisely, consider the check whether a regular product j (a flexible product k) can be sold given the current

state (\mathbf{c}, \mathbf{y}) . One could, of course, reduce the regular resources' remaining capacity to $\mathbf{c} - \mathbf{a}_j$ (increase the commitments to $\mathbf{y} + \mathbf{e}_k$), then apply FME, and finally check whether $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) \geq \mathbf{0}$ ($(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) \geq \mathbf{0}$) has a solution. However, it is not necessary to repeat FME so frequently throughout the booking horizon, which we will show in the following.

Proposition 2: Let $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ be an arbitrary state of DP-flex. Then, we have

- (a) $(\mathbf{c} - \mathbf{a}_j, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ if and only if $(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) \geq \mathbf{0} \forall j$
and
 $(\mathbf{c}, \mathbf{y} + \mathbf{e}_k) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ if and only if $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) \geq \mathbf{0} \forall k$
- (b) $(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) = (\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{a}}_j(\mathbf{A}, \mathcal{M})) \forall j$
and
 $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) = (\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{b}}_k(\mathbf{A}, \mathcal{M})) \forall k$

Proof: Expression (a) obviously follows from Proposition 1. Appendix B.1 provides the proof of expression (b).

Expression (a) implies that the decision whether a regular product j (a flexible product k) can be offered in state $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ need not be made by the original feasibility check (5)–(7), but can be equivalently made by using (11)–(12) instead. Expression (b) implies that, after a sale of a regular product j (of a flexible product k), FME need not be repeated. Instead, the regular and artificial resources capacity $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}))$ can simply be reduced by $(\mathbf{a}_j, \tilde{\mathbf{a}}_j(\mathbf{A}, \mathcal{M}))$ (by $(\mathbf{0}, \tilde{\mathbf{b}}_k(\mathbf{A}, \mathcal{M}))$), in order to obtain the transformed feasibility problem of the following state.

By exploiting the previous results, we can derive an alternative DP formulation. Consider an arbitrary instance of the transformed feasibility problem given by regular and artificial resources' remaining capacity \mathbf{c} and $\tilde{\mathbf{c}}$, respectively, as well as by the coefficients $\tilde{\mathbf{f}}_h$, $\tilde{\mathbf{b}}_k$, and $\tilde{\mathbf{a}}_j$. We define the DP's state space directly as $(\mathbf{c}, \tilde{\mathbf{c}})$. Let $V_t^{surr}(\mathbf{c}, \tilde{\mathbf{c}})$ denote the optimal expected revenue-to-go with t periods left, which can be computed recursively using the following Bellman equation (**DP-surr**):

$$V_t^{surr}(\mathbf{c}, \tilde{\mathbf{c}}) = \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot (r_j^{reg} + V_{t-1}(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}} - \tilde{\mathbf{a}}_j)) \right\}$$

$$\begin{aligned}
& + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \left(r_k^{flex} + V_{t-1}(\mathbf{c}, \tilde{\mathbf{c}} - \tilde{\mathbf{b}}_k) \right) \\
& + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}(\mathbf{c}, \tilde{\mathbf{c}}) \} \tag{13}
\end{aligned}$$

with the boundary conditions $V_t(\mathbf{c}, \tilde{\mathbf{c}}) = -\infty$ if $(\mathbf{c}, \tilde{\mathbf{c}}) \not\geq \mathbf{0}$ and $V_0(\mathbf{c}, \tilde{\mathbf{c}}) = 0 \forall (\mathbf{c}, \tilde{\mathbf{c}}) \geq \mathbf{0}$. Subsequently, we can formulate the following result:

Proposition 3: $V_t(\mathbf{c}, \mathbf{y}) = V_t^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}))$ for all t , (\mathbf{c}, \mathbf{y}) .

Proof: See Appendix B.2.

Propositions 2 and 3 imply that DP-flex (1) is indeed fully equivalent to DP-surr (13). The optimal expected value from the initial state is identical, and both DP formulations can be thought of as being carried out in parallel. More precisely, they allow the same decision options (due to Proposition 2 (a)) and virtually make the same decision (due to Proposition 2 (b) and Proposition 3). Thus, DP-surr can be used instead of DP-flex, we can drop the commitments from the state space, and, instead, track regular and artificial resources' capacity.

Accordingly, it suffices to apply FME only once to a given network, namely at the beginning of the booking horizon. In doing so, we fully retain the supply-side flexibility by postponing the assignment of flexible products. From a technical point of view, flexible products now correspond directly to resources. A flexible product is treated like a regular one; that is, the remaining capacity of artificial resources i is immediately reduced by \tilde{b}_{ik} , and the regular resources are left unchanged. In comparison, a regular product j requires—apart from the standard consumption \mathbf{a}_j —one unit of capacity of artificial resource i for every resource pooled in i and used by j ; that is, $\tilde{a}_{ij} = \sum_h \tilde{f}_{hi} \cdot a_{hj}$. Thus, DP-surr (13) has the same form as a standard revenue management problem without flexible products. This finally enables the application of standard solution approaches like DPD.

We call the network of products and resources underlying DP-surr the *surrogate* network of the original network underlying DP-flex. The surrogate networks are not only a technical output of FME, but are usually quite intuitive.

Example (cont'd): In the running example, the artificial resources are interpretable as upper limits on the total amount of the flexible product that can be sold. Obviously, the number of sales is restricted by the legs out of A ($\tilde{c}_1 = c_1 + c_3$) and into B ($\tilde{c}_3 = c_2 + c_4$). Additionally, remember that passengers travel either in the morning (legs 1 and 2) or in the afternoon (legs 3 and 4). If legs 1 and 4 are fully booked, no additional flexible product can be sold. This restriction is captured by $\tilde{c}_2 = c_1 + c_4$. Similarly, legs 2 and 3 may be the bottleneck ($\tilde{c}_4 = c_2 + c_3$). Together, these four artificial resources consider in an meaningful way all four combinations of bottlenecks for the number of sales that may occur and restrict sales of the flexible product to $\min\{c_1 + c_2\} + \min\{c_3 + c_4\}$.

Later, we also give an analogous interpretation of the example networks used in the numerical experiments (see Sections 5.2.2 and 5.3.2).

4.3 Network types and size of the surrogate networks

In the previous subsection, we have shown how an arbitrary revenue management problem with flexible products can be reformulated as an equivalent standard revenue management problem without flexible products. In this subsection, we focus on the size of the resulting surrogate network. This is an important issue, since FME can, in general, add a large number of constraints and there are examples where the number of added constraints is exponential in the problem size. Clearly, such examples also exist in the context of revenue management. When formulating the surrogate network, two or more of the m regular resources form an artificial resource, which a subset of the flexible products then uses. Thus, we have a maximum of $\tilde{m} = 2^m$ artificial resources. Examples that reach this upper bound can be easily constructed.

However, the networks (or, often, subnetworks), in which flexible products actually occur in practice, usually have a structure in which the number of artificial resources is mostly far less and stays polynomially bounded. It is important to observe that the number of artificial resources only depends on the structure of flexible products; that is, their alternatives \mathcal{M}_k . This number is completely independent of the regular products and also independent of flexible products' prices. That is, it does not increase if, besides an

existing flexible product k , a second flexible product k' is added with identical alternatives $\mathcal{M}_{k'} = \mathcal{M}_k$, but a different revenue and/or demand. Moreover, if arbitrary flexible products are deleted from a network, the number of artificial resources never increases. In the following, we thus focus on flexible products' structure and consider several network types whose flexible products are frequently used along with the resulting number of artificial resources.

Network type 1 (*fully flexible parallel resources*)

This network type consists of m parallel resources and a single flexible product that may be assigned to each of the m resources: $\mathbf{A} = [\mathbf{E}_{m \times m} | \cdot]$, $\mathcal{M}_1 = \{1, \dots, m\}$, where the columns to the right of $\mathbf{E}_{m \times m}$ denote the arbitrary resource consumption of regular products.

In practice, this network type occurs, for example, in the travel industry. Many tour operators offer travel roulette, which assigns customers to one of several similar hotels in their destination area.

In terms of the surrogate network, there is only one artificial resource that pools the capacity of all regular resources; that is, $\tilde{c}_1 = \sum_{h=1}^m c_h$. This resource is used by all products.

Proposition 4: In network type 1, the number of artificial resources is $\tilde{m} = 1$ (and thus constant in the number of regular resources m).

We show Proposition 4 by induction in Appendix C.1.

Network type 2 (*pairwise flexible parallel resources*)

This network type consists of m parallel resources and $m - 1$ flexible products. Flexible product k may be assigned to resource k or $k + 1$: $\mathbf{A} = [\mathbf{E}_{m \times m} | \cdot]$, $\mathcal{M}_k = \{k, k + 1\} \forall k$.

This network type arises, for example, in manufacturing. In a chain of factories, each factory k might be able to produce products $k - 1$ and k . When quantifying the benefit of flexibility, Jordan and Graves (1995) find that such a chain of factories yields nearly

the same output as a set of fully flexible factories, that can each produce every product. However, we leave out the last link in the chain here; there is no product that can be produced alternatively in factory m or 1. Another example is upgrading to the next higher resource in revenue management with parallel resources, for example, single-resource airline revenue management. Among others, Gallego and Stefanescu (2009) consider this upgrading “limited-cascading upgrading,” while Shumsky and Zhang (2009) consider it “single-step upgrading.” Moreover, the flexibility can also relate to time if, for example, guests’ alternative stays on a cruise ship are considered.

Regarding the surrogate network, for all $\underline{h} \in \{1, \dots, m-1\}$ and $\bar{h} \in \{\underline{h}+1, \dots, m\}$, there is an artificial resource $\tilde{c}_{\underline{h}\bar{h}} = \sum_{h=\underline{h}}^{\bar{h}} c_h$ that pools the capacity of the adjacent regular resources \underline{h} to \bar{h} . Artificial resource $\tilde{c}_{\underline{h}\bar{h}}$ is jointly used by flexible products \underline{h} to $\bar{h}-1$.

Proposition 5: In network type 2, the number of artificial resources is $\tilde{m} = \frac{(m-1) \cdot m}{2}$ (and thus polynomial in the number of regular resources m).

We show Proposition 5 by induction in Appendix C.2.

Network type 3 (*adjacent flexible parallel resources*)

This is a generalization of network type 2 and consists of m parallel resources, but $\frac{(m-1) \cdot m}{2}$ flexible products. More precisely, for all $\underline{h} \in \{1, \dots, m-1\}$ and $\bar{h} \in \{\underline{h}+1, \dots, m\}$, there is a flexible product that may be assigned to the adjacent resources \underline{h} to \bar{h} . Therefore, the flexible product k may be described in terms of the topologically first alternative \underline{h}_k and last alternative \bar{h}_k , and we have: $\mathbf{A} = [\mathbf{E}_{m \times m} | \cdot]$, $\mathcal{M}_k = \{h | \underline{h}_k \leq h \leq \bar{h}_k\} \forall k$.

In practice, this network type arises in generalizations of type 2, where, for example, a factory can produce more than two products; that is, some factories h are able to produce products $h-2$, $h-1$, and h . In single-resource revenue management, Gallego and Stefanescu (2009) have termed this generalization “full-cascading upgrading,” while Shumsky and Zhang (2009) call it “multi-step upgrading.”

We obtain the same set of artificial resources as in network type 2. However, the artificial resource $\tilde{c}_{\underline{h}\bar{h}}$ that pools capacity from \underline{h} to \bar{h} is now shared by all flexible products whose first, as well as last alternative, is between \underline{h} and \bar{h} ($\mathcal{M}_k \subseteq \{\underline{h}, \underline{h} + 1, \dots, \bar{h}\}$).

Proposition 6: In network type 3, the number of artificial resources is $\tilde{m} = \frac{(m-1) \cdot m}{2}$ (and thus polynomial in the number of regular resources m).

The structure of the proof is similar to the proof of Proposition 5 and omitted.

Network type 4 (*independent flexible block wise resources*)

This network type has a block structure, which consists of resource blocks $bl \in \mathcal{BL} = \{1, \dots, n^{block}\}$ with (w.l.o.g.) an identical number of m^{block} resources in each block, such that there are $m = m^{block} \cdot n^{block}$ resources in total. A flexible product k simultaneously uses resources from an arbitrary subset of blocks $\mathcal{BL}_k \subseteq \mathcal{BL}$. If k uses a block bl , the flexibility there is defined according to network type 3 (types 1 and 2 are also possible): $A^{bl} = [E_{m^{block} \times m^{block}} | \cdot]$, $\mathcal{M}_k^{bl} = \{h | \underline{h}_{k,bl} \leq h \leq \bar{h}_{k,bl}\} \forall k$. The key point here is that the blocks are independent in the sense that the final assignment of flexible product k to a regular resource in one block is independent of its assignment in other blocks; that is, an alternative $j \in \mathcal{M}_k$ combines arbitrary \mathcal{M}_k^{bl} for all $bl \in \mathcal{BL}_k$ and $|\mathcal{M}_k| = \prod_{bl \in \mathcal{BL}_k} |\mathcal{M}_k^{bl}|$.

In practice, this setting occurs, for example, in multi-stage production processes (see, e.g., Leachman and Carmon (1992)), in which the blocks correspond to the stages and the resources in a stage correspond to its machines. Another example is upgrading in network airline revenue management. Here, the blocks correspond to legs in the network and a ticket (i.e., a flexible product) encompasses one or more legs. On each leg, the passenger may be arbitrarily upgraded to a higher compartment (i.e., another resource in this block), for example, from economy to business class (see, e.g., Gönsch and Steinhardt (2015)).

In respect of the surrogate network, each resource block can be considered independently. In each block bl , we obtain $\frac{(m^{block}-1) \cdot m^{block}}{2}$ artificial resources. As in network types

2 and 3, each resource $\tilde{c}_{\underline{h}\bar{h}}$ pools the capacity of the adjacent regular resources \underline{h} to \bar{h} for all $\underline{h} \in \{1, \dots, m^{block} - 1\}$ and $\bar{h} \in \{\underline{h} + 1, \dots, m^{block}\}$.

Proposition 7: In network type 4, the total number of artificial resources is $\tilde{m} = n^{blocks} \cdot \frac{(m^{block-1}) \cdot m^{block}}{2} = \frac{m \cdot (m^{block-1})}{2}$ (and thus polynomial in the number of regular resources m).

The proof is similar to the proof of Proposition 5 and omitted.

Network type 5 (*dependent flexible block wise resources*)

Like network type 4, this network type has a block structure, which consists of resource blocks $bl \in \mathcal{BL} = \{1, \dots, n^{block}\}$. Again, a flexible product k simultaneously uses resources from an arbitrary subset of blocks $\mathcal{BL}_k \subseteq \mathcal{BL}$.

In contrast to network type 4, we now consider resource types. There are the same resource types $h \in \{1, \dots, n^{rt}\}$ in each resource block. Let the tuple (h, bl) refer to a resource of type h from a specific block bl , and let $c_{h,bl}$ denote its capacity, such that there are $m = n^{block} \cdot n^{rt}$ resources in total. The resource types follow a nested upgrade hierarchy. A higher index indicates a higher position in the hierarchy; that is, a more versatile resource.

The flexible product k is associated with resource type h_k ; that is, it can be assigned to h_k or upgraded to any $h > h_k$. Regarding block bl , we have: $\mathbf{A}^{bl} = [\mathbf{E}_{n^{rt} \times n^{rt}} | \cdot]$, $\mathcal{M}_{kh}^{bl} = \{h\} \forall k, h \geq h_k$. The important point here is that the assignment of product k must be the same $h \geq h_k$ for all blocks; that is, an alternative $j \in \mathcal{M}_k$ combines the \mathcal{M}_{kh}^{bl} for all $bl \in \mathcal{BL}_k$ and one $h \geq h_k$. Thus, we have $|\mathcal{M}_k| = n^{rt} - h_k + 1$.

The car rental industry is one of the most important users of such upgrades (see, e.g., Geraghty and Johnson (1997), Pachon et al. (2003), and Fink and Reiners (2006)). There, the blocks correspond to the days of the planning horizon and the resource types in a block correspond to different car types following a given upgrade hierarchy (e.g., economy, compact, and full-size car types). Another example is upgrading in the hotel industry, where the resource types correspond to different room types.

Regarding the surrogate network, we cannot consider the resource blocks independently anymore, because the assignments must be the same for all blocks. In the following, for all $h \in \{1, \dots, n^{rt} - 1\}$, we define $\mathcal{W}_h = \{(h, bl_h), (h + 1, bl_{h+1}), \dots, (n^{rt}, bl_{n^{rt}}) \mid bl_{h'} \in \mathcal{BL} \forall h' \in \{h, \dots, n^{rt}\}\}$. Each element $w \in \mathcal{W}_h$ refers to a set of resources with exactly one resource $(h', bl_{h'})$ of each type $h' \geq h$ from an arbitrary block $bl_{h'}$. Thus, w contains $n^{rt} - h + 1$ resources. As all combinations of blocks are considered, \mathcal{W}_h contains $(n^{block})^{n^{rt}-h+1}$ sets of resources. Now, there is one artificial resource \tilde{c}_{hw} for each $w \in \mathcal{W}_h$ and $h \in \{1, \dots, n^{rt} - 1\}$ that simply adds up the capacity of the resources in w :

$$\tilde{c}_{hw} = \sum_{(h', bl) \in w} c_{h', bl} \quad \forall w \in \mathcal{W}_h, h \in \{1, \dots, n^{rt} - 1\} \quad (14)$$

Let \mathcal{BL}^w denote the set of resource blocks from which resources are contained in $w \in \mathcal{W}_h$. The artificial resource \tilde{c}_{hw} is shared by all products with $h_k \geq h$ and $\mathcal{BL}_k \supseteq \mathcal{BL}^w$.

Proposition 8: In network type 5, the total number of artificial resources is $\tilde{m} = \sum_{k=2}^{n^{rt}} (n^{block})^k$ (and thus polynomial in the number of blocks and exponential in the number of resource types).

The structure of the proof is similar to the proof of Proposition 5 and is omitted.

Please note that the number of resource types is relatively small and constant in most practical applications. For example, in the car rental industry, there are often three to five car types in the upgrade hierarchy. In contrast, the number of resource blocks considered varies across rental stations and is often subject to the individual decision maker. Thus, also in this setting, the problem stays polynomially bounded in the relevant, potentially scalable problem parameters, that is, the resource blocks.

5 Computational experiments

In this section, we evaluate the revenue performance of the surrogate approach from Section 4. We use two airline networks introduced by Liu and van Ryzin (2008), which became de facto standard test instances for choice-based revenue management (see, e.g.,

Miranda Bront et al. (2009); Meissner and Strauss (2012)). We describe the experimental setup in Section 5.1, and we evaluate the approaches' revenue performance in detail in Sections 5.2 and 5.3, separately for the two networks under consideration.

5.1 Experimental setup

We summarize the implemented revenue management methods in Section 5.1.1. Please note that the technical details are provided in Appendix D. Furthermore, we describe the customer choice behavior in Section 5.1.2 and explain the consideration of forecast uncertainty in Section 5.1.3.

5.1.1 Implemented revenue management methods

Our main method is *DPD-surr*, which implements the surrogate approach described in the previous section. In this method, the surrogate reformulation is solved with the DPD approach of Liu and van Ryzin (2008). Details can be found in Appendix D.4. As benchmarks, we implemented the two methods *DPD-ah* and *CDLP-surr* as well as an upper bound (*UB*) on the optimal expected revenue of DP-flex (1):

- *DPD-ah* is a DPD approach that forgoes flexibility and immediately assigns flexible products (ad hoc) after sale (see Appendix D.2). Several studies report a good revenue performance of this approach in settings with independent demand (see Section 2). We incorporated this ad hoc assignment into the DPD approach of Liu and van Ryzin (2008).
- *CDLP-surr* refers to the optimal primal solution of the corresponding CDLP formulation (D.1.8)–(D.1.12) that gives us the time a set S should be offered during the booking horizon (see Appendix D.3). This is in line with a benchmark used by Liu and van Ryzin (2008).
- *UB* is the upper bound obtained from the optimal objective value of the CDLP (see Appendix D.1). This value can be obtained by either solving the CDLP model with flexible products (CDLP-flex (D.1.1)–(D.1.6); see Gallego et al. (2004)) or by using the surrogate reformulation in the standard CDLP formulation without flexible

products (CDLP-surr (D.1.8)–(D.1.12); see, e.g., Liu and van Ryzin (2008) and Miranda Bront et al. (2009)).

All algorithms were implemented in MATLAB (Version 8, Release R2013a). Linear programs were solved by the function `linprog` from the Optimization Toolbox. We use Monte Carlo simulation to evaluate the described methods and report values averaged over 200 customer streams for each problem instance.

5.1.2 Customer choice behavior

We assume the same choice behavior as Liu and van Ryzin (2008). Therefore, the choice model and the notation required to describe our computational experiments are only summarized in brief. Each customer belongs to a segment $l \in \mathcal{L}$, and customers from l are only interested in a subset of the entire product set, namely their consideration set C_l . Furthermore, the consideration sets are disjoint for customers belonging to different segments. With probability λ_l , a customer from segment l arrives. Her segment-specific purchase probabilities, that is, $P_{lj}^{reg}(S)$ for regular product j , $P_{lk}^{flex}(S)$ for flexible product k , and $P_{l0}(S)$ for the no-purchase alternative, are given by the standard multinomial logit model. They are computed using her product-specific preference weights, denoted by the parameters v_{lj}^{reg} , v_{lk}^{flex} , and v_{l0} for regular product j , flexible product k , and the no-purchase alternative, respectively. Then, for this choice model, the purchase probability is computed by

$$P_{lj}^{reg}(S) = \frac{v_{lj}^{reg}}{\sum_{j \in C_l \cap S} v_{lj}^{reg} + \sum_{k \in C_l \cap S} v_{lk}^{flex} + v_{l0}} \quad (15)$$

for a regular product. For a flexible product and the no-purchase alternative, only the numerator changes. Please note that, because of the assumption of a multinomial logit model and disjoint consideration sets, the large number of (column generation) subproblems arising in *DPD-surr*, *DPD-ah*, *CDLP-surr*, and *UB* (i.e., the problems determining the offer set) can be solved efficiently by a simple ranking procedure (see Liu and van Ryzin (2008)).

Regarding the segment probabilities' temporal distribution, we consider two arrival patterns. The first one models *time-homogenous* demand. In the second arrival pattern, called *mixed*, we consider that low-value demand tends to arrive earlier. This pattern is obtained by assuming that 50% of demand is time-homogenous, and 50% arrives according to the classical low-before-high assumption. We straightforwardly adapt our models by introducing time-dependent arrival probabilities in the DPD approaches and in the CDLP approximations used therein.

5.1.3 Forecast uncertainty

To incorporate forecast uncertainty, we implemented stochastic forecast errors as studied in Petrick et al. (2012) to disturb the regular products' preference weights. The forecast errors are itinerary-based. All forecasted preference weights concerning a specific itinerary are disturbed by the same factor. Therefore, a uniformly distributed random number $\hat{\delta} \in \mathcal{U}(-\delta, +\delta)$ is drawn within every simulation run for each of the regular products' itineraries, and the corresponding preference weights are multiplied by the factor $(1 + \hat{\delta})$. The size of the error is controlled by the error bound $\delta \in [0,1]$.

5.2 Network 1: Parallel flights

In the following sections, we explain how we modified the first example from Liu and van Ryzin (2008) to include flexible products (Section 5.2.1) and interpret the corresponding surrogate reformulation (Section 5.2.2). We then evaluate the approaches' revenue performance in detail (Section 5.2.3).

5.2.1 Network description

Network 1 consists of three parallel legs with capacity $\mathbf{c} = (30, 50, 40)^T$ that can be thought of as flights on the same route at different times of day. On each leg, the firm offers a high fare class regular product (products 1–3) and a low fare class regular product (products 4–6). The prices are given by $\mathbf{r}^{reg} = (800, 1000, 600, 400, 500, 300)^T$. In addition, we consider a flexible product (product

f) which can be sold at a price of $r_f^{flex} = 240$. The flexible product guarantees transportation on one of the three legs.

There is a high fare class customer segment H (consideration set $C_H = \{1,2,3\}$), a low fare class segment L (consideration set $C_L = \{4,5,6\}$) and a flexible segment F (consideration set $C_F = \{f\}$). The preference vectors are given by $v_H = (5,10,1)^T$, $v_L = (5,1,10)^T$, and $v_F = (10)$.

As usual in revenue management experiments, we tested different network loads. We varied the scarcity of capacity using a capacity factor $\alpha \in \{0.4, 0.5, \dots, 1.2\}$. Different customer attitudes were captured by four no-purchase preference vectors $v_0 = (v_{H0}, v_{L0})^T$, that is, $(0.01, 0.01)^T$, $(1, 5)^T$, $(5, 10)^T$, and $(10, 20)^T$. The flexible customer segment's preference weight for the no-purchase alternative is 0.01, and the number of periods is set to 300.

In the *time-homogenous* arrival pattern, customers of segments H , L , and F arrive with probabilities $\lambda_H = 0.2$, $\lambda_L = 0.3$, and $\lambda_F = 0.1$, respectively. Accordingly, in the *mixed* arrival pattern, the probabilities $(\lambda_H, \lambda_L, \lambda_F)^T$ are $(0.1, 0.15, 0.35)^T$, $(0.1, 0.45, 0.05)^T$, and $(0.4, 0.15, 0.05)^T$ in periods 300–251, 250–101, and 100–1, respectively.

5.2.2 Surrogate reformulation

The surrogate network consists of four resources: the three regular resources and one artificial resource with capacity $c_1 + c_2 + c_3$. It can readily be interpreted. The artificial resource pools the capacity of all resources that can potentially be used to fulfill the flexible product. More precisely, it represents the maximum amount of flexible and regular products that can be sold, because they all jointly use the capacity of the three legs. Accordingly, if a product is actually sold, the artificial resource's capacity is reduced. A flexible product needs capacity on this artificial resource alone, because the number of flexible product sales is constrained only by the joint capacity of the three legs. By contrast, a regular product requires one unit of capacity on 'its' regular resource and one unit of capacity on the artificial resource. The consumption of the regular resource re-

flects that one seat fewer is now available on this leg. The consumption of the artificial resource reflects that the seat can be used neither for a flexible nor a regular product.

5.2.3 Performance evaluation

Figure 2 shows the average revenues of *DPD-surr*, *DPD-ah*, and *CDLP-surr* relative to *UB* in all scenarios subject to the capacity factor α . Each column relates to a specific no-purchase preference vector, and each row represents one of the two arrival patterns (*time-homogenous* or *mixed*). Forecast errors are not considered here.

In general, all three methods' revenue performance is rather good, as expected from the literature on DPD without flexible products (see, e.g., Miranda Bront et al. (2009)). There seems to be no major impact of the arrival pattern. *DPD-ah* usually yields 94%–98% of *UB* and *CDLP-surr* yields 96%–98%. *DPD-surr* attains even higher revenues of 97%–99% of *UB*. As usual, revenue management is relatively easy for extreme network load factors. If capacity is very scarce ($\alpha = 0.4$ and $(v_{0H}, v_{0L})^T = (0.01, 0.01)^T$), only the high fare class products are offered. Similarly, revenue management becomes more or less obsolete when all products are offered in case of ample capacity ($\alpha \geq 1.0$), and all methods yield revenues close to *UB*. But for intermediate capacity, where revenue management is most relevant, considerable differences can be observed. Here, *DPD-surr* shows a very stable revenue performance, whereas *DPD-ah* yields considerably lower revenues in many cases. This is most obvious for preference weights of $(v_{0H}, v_{0L})^T = (0.01, 0.01)^T$. In both the *time-homogeneous* and the *mixed* arrival patterns, *DPD-ah*'s revenue falls to under 90% at $\alpha = 0.8$, whereas *DPD-surr* still attains about 98% and *CDLP-surr* remains at 96%.

Next, we focus on the relative performance of the two DPD methods and consider forecast errors. Figure 3 shows the revenue gain of *DPD-surr* over *DPD-ah*, subject to the upper error bound δ on the forecast uncertainty. To keep the figure simple, we only depict the most relevant capacity factors ($\alpha \in \{0.5, 0.6, 0.7, 0.8\}$). Furthermore, we tested whether these revenue gains are significant at the 99% level of confidence. We calculated the revenue difference together with the empirical standard deviation on a per-stream

basis and conducted a standard paired t-test. If the 99% confidence interval of the revenue difference does not include zero, the gain is significant. For reasons of clarity and because all confidence intervals are similar in size, we only included error bars for the top and bottom lines in the plots of Figure 3.

For the original setting without forecast errors ($\delta = 0$), we observe revenue gains of around 1% and 2% in the majority of cases. In general, the gains increase with higher forecast uncertainty. The higher the δ , the more important it is to use flexible products to mitigate demand uncertainty. Obviously, *DPD-surr* can benefit considerably from retaining full flexibility of the requests already accepted. This is in line with an observation from Petrick et al. (2010) who obtained similar results regarding linear programming-based heuristics that retained flexibilities to varying degrees.

At first glance, it seems strange that there is almost no influence of the forecast error in network 1 for preference weights of $(v_{0H}, v_{0L})^T = (0.01, 0.01)^T$. This is due to the special demand structure in this standard setting where the probability for the no-purchase alternative is almost zero, as long as a product can be bought ($P_{l0}(S) \approx 0 \forall S \cap C_l \neq \emptyset$). The customers' preference weights, which are disturbed by the forecast error, essentially do not influence whether a customer buys, they only influence which product she buys. However, correctly anticipating this decision is not important, because the products in a customer's consideration set have similar revenues, and if a leg is fully booked, only the other legs' products are offered and bought with probability one.




Figure 2: Average revenues of *DPD-surr*, *DPD-ah*, and *CDLP-surr* relative to *UB* in network 1

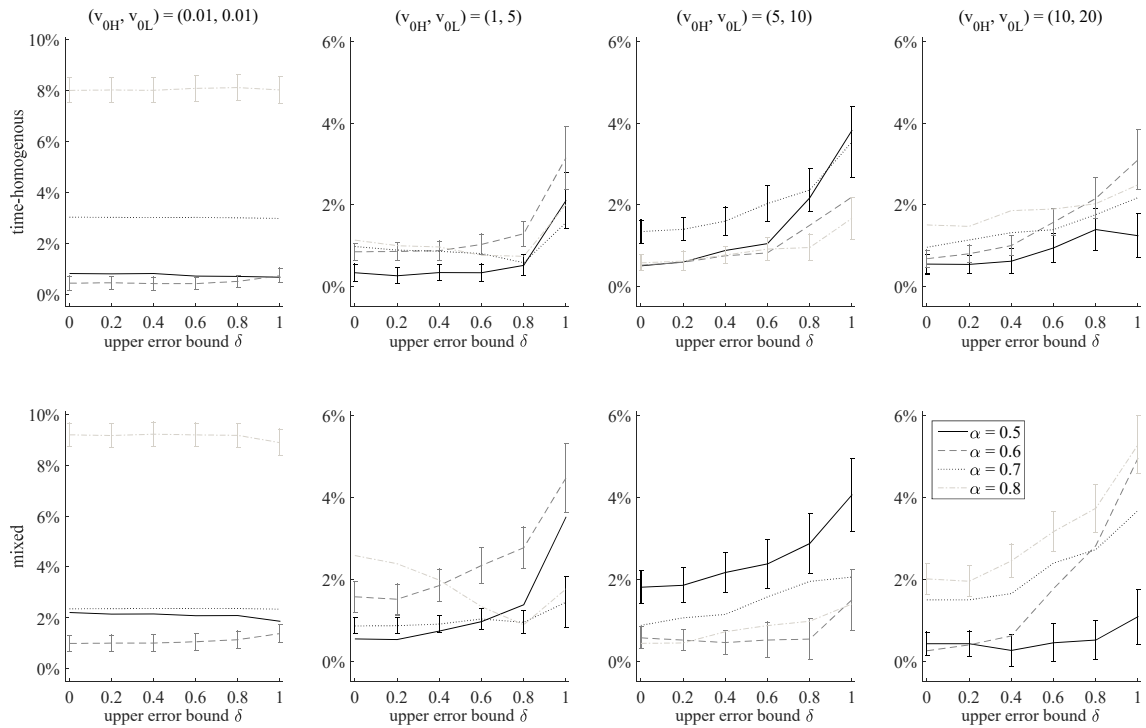


Figure 3: Revenue gains of *DPD-surr* over *DPD-ah* *UB* in network 1

5.3 Network 2: Small hub-and-spoke network

Again, we first describe the specific product and demand data of the second problem instance (Section 5.3.1) and interpret the corresponding surrogate reformulation (Section 5.3.2). We then turn to the computational results (Section 5.3.3).

5.3.1 Network description

Network 2 consists of seven flight legs connecting the four cities A, B, C, and H (see Figure 4). There are 11 itineraries, and on each itinerary, the airline offers a high fare class and a low fare class product. Details on prices and capacity consumption of these 22 regular products are identical to Liu and van Ryzin (2008) and provided in Table E.1 in Appendix E. In addition, we consider five flexible products (products $f1$ – $f5$). The first one offers transportation from A to B, either on leg 1, legs 2 and 4, or legs 3 and 5. The second flexible product is from A to C on either (2, 6) or (3, 7). The last three flexible products guarantee short-haul transportation on one of the two possible legs from A to H, H to B, and H to C. The prices of these five products are given by $\mathbf{r}^{flex} = (240, 280, 160, 120, 200)^T$. We consider 15 customer segments: one high fare class and one low fare class segment interested in regular products for each of the origin-destination pairs AB, AH, HB, HC, and AC, and one segment for each of the five flexible products. Details on consideration sets, preference vectors, and segment probabilities can be found in Table E.2 in Appendix E. Different customer attitudes are again captured by the four values of the no-purchase preference vector already used in network 1. The flexible product segments' weights for the no-purchase alternative are fixed at 0.01. Analogously to network 1, we consider the two arrival patterns *time-homogenous* and *mixed* as well as a capacity factor.

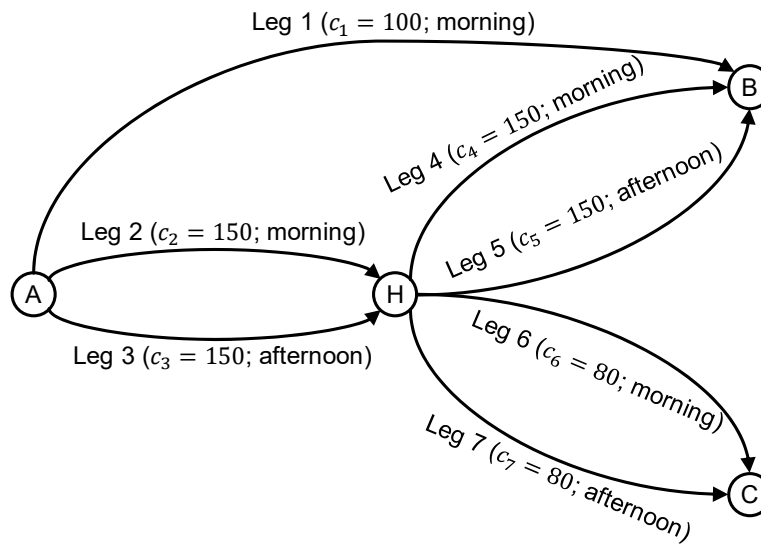


Figure 4: Small hub-and-spoke network (Network 2; see Liu and van Ryzin (2008))

5.3.2 Surrogate reformulation

The surrogate network comprises 11 artificial resources. The artificial resources and flexible products' capacity consumption are shown in Table 3.

Regular products require one unit of capacity of the corresponding regular resource(s) (see Table E.1) and one unit of capacity of each artificial resource containing the regular resource.

Again, the artificial resources are interpretable. For example, sales of f_4 (H to B) are limited by the joint capacity of legs 4 and 5. This restriction is captured by the artificial resource $\tilde{c}_5 = c_4 + c_5$. Note that the restriction imposed by artificial resource $\tilde{c}_9 = c_1 + c_4 + c_5$ is obviously weaker and never limiting for f_4 sales, but it is necessary to capture an interaction with f_1 , which will be described later.

Product f_1 (A to B) is a bit more tedious. Similar to the running example from Section 4, the number of sales is restricted by the joint capacity of the legs out of A ($\tilde{c}_{11} = c_1 + c_2 + c_3$) and into B ($\tilde{c}_9 = c_1 + c_4 + c_5$). Furthermore, legs 2 and 5 ($\tilde{c}_1 = c_1 + c_2 + c_5$) or legs 3 and 4 ($\tilde{c}_2 = c_1 + c_3 + c_4$) may be the bottleneck. Together, these four artificial resources restrict sales of f_1 to $c_1 + \min\{c_2 + c_4\} + \min\{c_3 + c_5\}$.

Besides the restrictions on sales for individual flexible products, the capacity shared by multiple flexible products has to be taken into consideration. Artificial resource 9 is a simple example: Capacity on legs 4 and 5 used by $f4$ customers cannot be used by $f1$ customers. Thus, this artificial resource – which was derived above as an individual restriction for $f1$ – is in fact not only used by $f1$ but also by $f4$. Furthermore, there can also be additional artificial resources that are only required because of such interactions between flexible products. Artificial resource 6 ($\tilde{c}_6 = c_1 + c_2 + c_5 + c_7$) is an example of this: It restricts joint sales of $f1$ and $f2$, because customers going from A to B ($f1$) or A to C ($f2$) either leave in the morning (legs 1 and 2) or arrive in the afternoon (legs 5 and 7). Note that this interaction between $f1$ and $f2$ is not captured by artificial resources 1, 2, 9, and 11 described above, because the routing of $f2$ may also be restricted by capacity on legs 6 and 7. An analogous example is $\tilde{c}_7 = c_1 + c_3 + c_4 + c_6$.

Artificial resource		Flexible product				
Index i	Capacity \tilde{c}_i	$f1$	$f2$	$f3$	$f4$	$f5$
1	$c_1 + c_2 + c_5$	X				
2	$c_1 + c_3 + c_4$	X				
3	$c_2 + c_7$		X			
4	$c_3 + c_6$		X			
5	$c_4 + c_5$				X	
6	$c_1 + c_2 + c_5 + c_7$	X	X			
7	$c_1 + c_3 + c_4 + c_6$	X	X			
8	$c_2 + c_3$		X	X		
9	$c_1 + c_4 + c_5$	X			X	
10	$c_6 + c_7$		X			X
11	$c_1 + c_2 + c_3$	X	X	X		

Table 3: Artificial resources of small hub-and-spoke network

5.3.3 Performance evaluation

Analogously to network 1, Figure 5 and Figure 6 show the average revenues of *DPD-surr*, *DPD-ah*, and *CDLP-surr*, as well as the revenue gain of *DPD-surr* over *DPD-ah* in network 2.

Compared with network 1, average revenues in the standard setting without forecast errors (Figure 5) are slightly higher in network 2 for both *DPD-ah* (97.5%–99%) and *DPD-surr* (99%–100%). Again, the performance of *DPD-surr* is more stable without any outliers, while *DPD-ah*'s revenue often falls below 98% of *UB* for capacity factors of $\alpha = 0.8$. By contrast, *CDLP-surr* performs considerably worse with many revenues around 93% and only a few values exceeding 96%.

For $(v_{0H}, v_{0L})^T = (0.01, 0.01)^T$, the revenue gain (Figure 6) of *DPD-surr* is considerably smaller at around 1%–2%. With the other three preference weights of $(v_{0H}, v_{0L})^T$ considered, the gain is more or less the same as in network 1. However, especially for $(v_{0H}, v_{0L})^T = (10, 20)^T$ and time-homogeneous demand, considerably higher gains are observed.

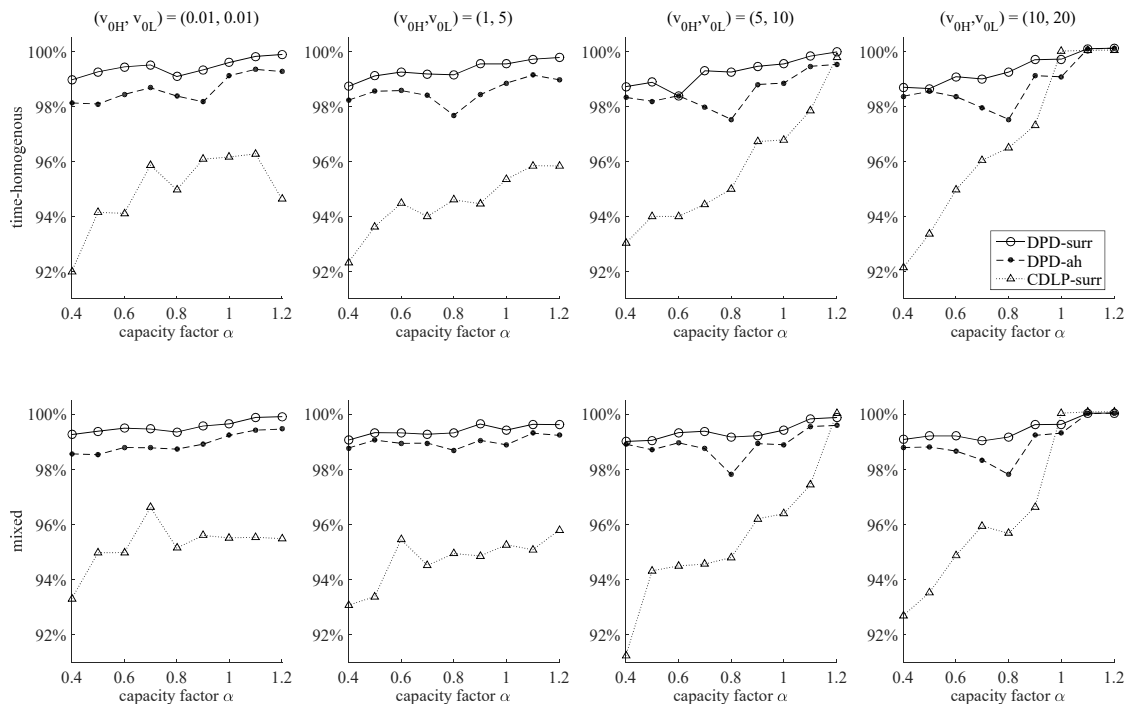


Figure 5: Average revenues of *DPD-surr*, *DPD-ah*, and *CDLP-surr* relative to *UB* in network 2

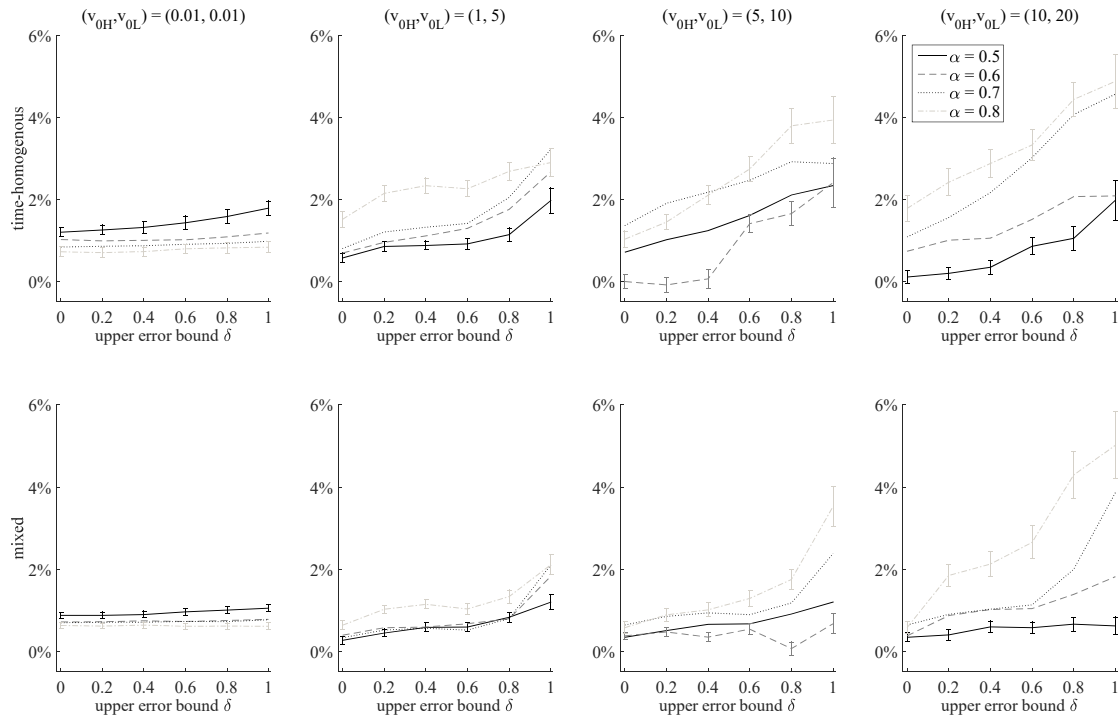


Figure 6: Revenue gains of *DPD-surr* over *DPD-ah UB* in network 2

6 Discussion and future research

Several managerial implications follow from our work. Most importantly, the inclusion of flexible products no longer excludes the use of standard dynamic programming techniques. We presented a novel generic way to overcome the commitment-based state space and the feasibility problem inherent in network revenue management problems regarding flexible products. In the surrogate approach, the problem is reformulated by applying FME to the feasibility problem, and an equivalent standard revenue management problem is obtained. This allows the direct use of standard DPD. Moreover, it allows the continued use of arbitrary methods and existing software systems, albeit with modified input data.

In a large number of numerical experiments, we compared the approach with a benchmark approach adapted from the literature that forgoes flexibility and obtains a resource-based state space by immediately assigning flexible products (ad hoc) after sale. The surrogate approach consistently obtains the highest revenues, which are close to the theoretical upper bound. In test instances with intermediate capacity, this approach in-

creases revenues by up to 8% compared to the ad hoc approach. Moreover, revenue gains increase when forecast errors are considered. Thus, we think the surrogate approach should be the first choice when incorporating flexible products into revenue management.

Moreover, the difference between the revenue of the surrogate approach and that obtained with the ad hoc approach can also be roughly interpreted as the supply-side benefit of offering flexible products instead of opaque products. Flexible products should be offered if this benefit outweighs their demand-side disadvantages (customers usually prefer an opaque product where they are immediately informed of what they get). However, there is no clear advice here. The difference is marginal in some cases (extreme capacity situations, low forecast errors) and considerable in others (intermediate capacity, medium to high forecast errors).

We think that our results are promising, and we encourage future work on this topic. First, research could focus on problem instances where our transformed feasibility problem is heuristic. In this respect, a starting point might be projection for integer problems (see, e.g., Williams and Hooker (2014)). Second, our results indicate that heuristics restricting flexible products' flexibility can also yield a good revenue performance. Consequently, we think it is promising to develop approximate dynamic programming techniques tailored to flexible products that retain more flexibility than the ad hoc approach. For example, the linear programming approach for approximate dynamic programming (see, e.g., Adelman (2007) for the traditional revenue management setting; Tong and Topaloglu (2014), as well as Vossen and Zhang (2015a) for refinements) could be extended to flexible products.

Appendix

A Proof of Proposition 1

Proposition 1: Given Condition 1 holds, $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})}$ (that is, the feasibility problem (2)–(4) has a solution) if and only if $(\mathbf{c}, \tilde{\mathbf{c}}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y})) \geq \mathbf{0}$ (that is, (11)–(12) has a solution).

Proof: First, (2)–(4) is reformulated as (8) using Condition 1. This is transformed into (11)–(12) by projecting out one distribution variable after the other, using Algorithm 1. To show that the whole algorithm keeps equivalence, it is sufficient to show that the inequality system before an iteration implies the inequality system after the iteration and vice versa. In the following, we show this for an arbitrary iteration d . Please note that, due to the construction of the algorithm, d always refers to the first column of the current **LHS**, which contains at least one nonzero coefficient.

W.l.o.g., assume that the coefficients in the first column are 1, -1 , or 0. The system before the iteration starts is given by

$$1 \cdot y_d + \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{pos,col} \cdot y_{col} \leq \mathbf{RHS}_{pos} \text{ for all } pos \in Pos \quad (\text{A.1})$$

$$(-1) \cdot y_d + \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{neg,col} \cdot y_{col} \leq \mathbf{RHS}_{neg} \text{ for all } neg \in Neg \quad (\text{A.2})$$

$$\sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{null,col} \cdot y_{col} \leq \mathbf{RHS}_{null} \text{ for all } null \in Null \quad (\text{A.3})$$

and the system after the iteration by

$$\sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{null,col} \cdot y_{col} \leq \mathbf{RHS}_{null} \text{ for all } null \in Null \quad (\text{A.4})$$

$$\sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} (\text{lhs}_{pos,col} + \text{lhs}_{neg,col}) \cdot y_{col} \leq \mathbf{RHS}_{pos} + \mathbf{RHS}_{neg} \quad \text{for all } (pos, neg) \in Pos \times Neg. \quad (\text{A.5})$$

Now, note that (A.1) and (A.2) imply

$$-\mathbf{RHS}_{neg} + \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{neg,col} \cdot y_{col} \leq y_d \leq \mathbf{RHS}_{pos} - \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{pos,col} \cdot y_{col} \quad \text{for all } (pos, neg) \in Pos \times Neg, \quad (\text{A.6})$$

which itself implies

$$-\mathbf{RHS}_{neg} + \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{neg,col} \cdot y_{col} \leq \mathbf{RHS}_{pos} - \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{pos,col} \cdot y_{col}$$

for all $(pos, neg) \in Pos \times Neg$. (A.7)

The proof that (A.1)–(A.3) implies (A.4)–(A.5) is particularly easy. Consider an arbitrary solution $(y_d, \dots, y_{|\sum_k \mathcal{M}_k|})^T$ from (A.1)–(A.3). Then, (A.4) has a solution, because it is equivalent to (A.3). Moreover, because (A.1) and (A.2) hold, (A.7) holds and, therefore, also (A.5), which is (A.7) slightly reformulated.

Now, in order to show that (A.4)–(A.5) implies (A.1)–(A.3), consider an arbitrary solution $(y_{d+1}, \dots, y_{|\sum_k \mathcal{M}_k|})^T$ from (A.4)–(A.5). (A.3) has a solution, because it is equivalent to (A.4). Moreover, because (A.5) and its reformulation (A.7) have a solution, we can construct a y_d for which $y_d \geq \max_{neg \in Neg} \left\{ -\mathbf{RHS}_{neg} + \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{neg,col} \cdot y_{col} \right\}$ and $y_d \leq \min_{pos \in Pos} \left\{ \mathbf{RHS}_{pos} - \sum_{col=d+1}^{|\sum_k \mathcal{M}_k|} \text{lhs}_{pos,col} \cdot y_{col} \right\}$ holds. Therefore (A.2) and (A.1) hold, too. \square

B Equivalence of DP-flex and DP-surr

B.1 Proof of Proposition 2

Proposition 2: Let $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(A, \mathcal{M})}$ be an arbitrary state of DP-flex. Then, we have

(a) $(\mathbf{c} - \mathbf{a}_j, \mathbf{y}) \in \mathcal{Z}_{(A, \mathcal{M})}$ if and only if $(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) \geq \mathbf{0} \forall j$
and

$$(\mathbf{c}, \mathbf{y} + \mathbf{e}_k) \in \mathcal{Z}_{(A, \mathcal{M})} \text{ if and only if } (\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) \geq \mathbf{0} \forall k$$

(b) $(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) = (\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{a}}_j(A, \mathcal{M})) \forall j$

and

$$(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) = (\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{b}}_k(A, \mathcal{M})) \forall k$$

Proof: Expression (a) obviously follows from Proposition 1, because we can simply define the state $(\mathbf{c} - \mathbf{a}_j, \mathbf{y}) := (\mathbf{c}', \mathbf{y})$ (and $(\mathbf{c}, \mathbf{y} + \mathbf{e}_k) := (\mathbf{c}, \mathbf{y}')$) and apply Proposition 1 to the so-defined state.

To show expression (b), first consider the upper case; that is, the sale of a regular product j . Regarding the first term in both brackets $(\mathbf{c} - \mathbf{a}_j$, i.e., the regular resources), the

equality is trivial. Regarding the second term in the brackets (i.e., the artificial resources), we have:

$$\begin{aligned}
\tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y}) &= \\
&= \sum_h \tilde{f}_h(\mathbf{A}, \mathcal{M}) \cdot (c_h - a_{hj}) - \sum_k \tilde{b}_k(\mathbf{A}, \mathcal{M}) \cdot y_k = \\
&= \sum_h \tilde{f}_h(\mathbf{A}, \mathcal{M}) \cdot c_h - \sum_k \tilde{b}_k(\mathbf{A}, \mathcal{M}) \cdot y_k - \sum_h \tilde{f}_h(\mathbf{A}, \mathcal{M}) \cdot a_{hj} = \\
&= \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{a}_j(\mathbf{A}, \mathcal{M})
\end{aligned}$$

The first equality simply follows from the definition of the function $\tilde{c}(\cdot)$ with reduced capacity, the second is algebra, and the third uses the definitions of $\tilde{c}(\cdot)$ and $\tilde{a}_j(\cdot)$.

Next, consider the lower case of Proposition 2 (b); that is, the sale of a specific flexible product k (in the summations, flexible products are denoted as k' in the following). Similarly to the considerations above, we only have to consider the second term in brackets:

$$\begin{aligned}
\tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k) &= \\
&= \sum_h \tilde{f}_h(\mathbf{A}, \mathcal{M}) \cdot c_h - \sum_{k' \neq k} \tilde{b}_{k'}(\mathbf{A}, \mathcal{M}) \cdot y_{k'} - \tilde{b}_k(\mathbf{A}, \mathcal{M}) \cdot (y_k + 1) = \\
&= \sum_h \tilde{f}_h(\mathbf{A}, \mathcal{M}) \cdot c_h - \sum_{k'} \tilde{b}_{k'}(\mathbf{A}, \mathcal{M}) \cdot y_{k'} - \tilde{b}_k(\mathbf{A}, \mathcal{M}) \cdot 1 = \\
&= \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{b}_k(\mathbf{A}, \mathcal{M})
\end{aligned}$$

The first and the third equality follow from the definition of $\tilde{c}(\cdot)$ with increased commitments and the second is algebra. \square

B.2 Proof of Proposition 3

Proposition 3: $V_t(\mathbf{c}, \mathbf{y}) = V_t^{surr}(\mathbf{c}, \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y}))$ for all t , (\mathbf{c}, \mathbf{y}) .

Proof: The equality is shown by induction over t . It holds for $t = 0$, because, from the boundary conditions, we have $V_0(\mathbf{c}, \mathbf{y}) = V_0^{surr}(\mathbf{c}, \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y})) = 0$ for $(\mathbf{c}, \mathbf{y}) \in \mathcal{Z}_{(\mathbf{A}, \mathcal{M})} \Leftrightarrow (\mathbf{c}, \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y})) \geq \mathbf{0}$ and $V_0(\mathbf{c}, \mathbf{y}) = V_0(\mathbf{c}, \tilde{c}(\mathbf{A}, \mathcal{M}, \mathbf{c}, \mathbf{y})) = -\infty$ otherwise, where the equivalence is Proposition 1.

Now, assume that the result holds for $t - 1$. In respect of t , two cases have to be distinguished again. If the boundary condition $(\mathbf{c}, \mathbf{y}) \notin \mathcal{Z}_{(A, \mathcal{M})} \Leftrightarrow (\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y})) \not\geq \mathbf{0}$ applies, we have $V_t(\mathbf{c}, \mathbf{y}) = V_t^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y})) = -\infty$. Otherwise, we have:

$$\begin{aligned}
V_t(\mathbf{c}, \mathbf{y}) &= \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} + V_{t-1}(\mathbf{c} - \mathbf{a}_j, \mathbf{y}) \right) \right. \\
&\quad \left. + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \left(r_k^{flex} + V_{t-1}(\mathbf{c}, \mathbf{y} + \mathbf{e}_k) \right) \right. \\
&\quad \left. + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}(\mathbf{c}, \mathbf{y}) \right\} \\
&= \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} + V_{t-1}^{surr}(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c} - \mathbf{a}_j, \mathbf{y})) \right) \right. \\
&\quad \left. + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \left(r_k^{flex} + V_{t-1}^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y} + \mathbf{e}_k)) \right) \right. \\
&\quad \left. + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y})) \right\} \\
&= \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} + V_{t-1}^{surr}(\mathbf{c} - \mathbf{a}_j, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{a}}_j(A, \mathcal{M})) \right) \right. \\
&\quad \left. + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \left(r_k^{flex} + V_{t-1}^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y}) - \tilde{\mathbf{b}}_k(A, \mathcal{M})) \right) \right. \\
&\quad \left. + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y})) \right\} \\
&= V_t^{surr}(\mathbf{c}, \tilde{\mathbf{c}}(A, \mathcal{M}, \mathbf{c}, \mathbf{y}))
\end{aligned}$$

The first equality is simply the definition of DP-flex (1), the second uses the induction hypothesis, the third equality follow from Proposition 2 (b), and the fourth is the definition of DP-surr (13). \square

C Derivation of artificial resources for network types 1 and 2

C.1 Proof of Proposition 4

Proposition 4: In network type 1, the number of artificial resources is $\tilde{m} = 1$ (and thus constant in the number of regular resources m).

Proof: Network type 1 consists of m parallel resources and one flexible product that may be assigned to the m resources. Thus, the feasibility problem (5)–(7) is given by

$$y_{1h} \leq c_h \quad \forall h = 1, \dots, m \quad (\text{C.1.1})$$

$$\sum_{h=1}^m -y_{1h} \leq -y_1 \quad (\text{C.1.2})$$

$$-y_{1h} \leq 0 \quad \forall h = 1, \dots, m \quad (\text{C.1.3})$$

Throughout the proof, we refer to the FME-steps as given by Algorithm 1. W.l.o.g., we assume that the elimination of y_{1h} by Algorithm 1 is done iteratively in increasing order of h . We index the iterations of Algorithm 1 with $i = 1, \dots, m + 1$, each referring to the feasibility problem before eliminating the distribution variable y_{1i} (and after eliminating $y_{1,i-1}$). Please note that iteration $i = 1$ corresponds to the initial feasibility problem and that the dummy iteration $i = m + 1$ gives us the feasibility problem after eliminating all the distribution variables.

Now, by induction over i , we show the following:

Induction hypothesis: The feasibility problem in iteration i is given by

$$y_{1h} \leq c_h \quad \forall h = i, \dots, m \quad (\text{C.1.4})$$

$$-y_{1h} \leq 0 \quad \forall h = i, \dots, m \quad (\text{C.1.5})$$

$$0 \leq c_h \quad \forall h = 1, \dots, i - 1 \quad (\text{C.1.6})$$

$$-\sum_{h=i}^m y_{1h} \leq \sum_{h=1}^{i-1} c_h - y_1 \quad (\text{C.1.7})$$

Induction basis: The induction hypothesis holds for $i = 1$, because (C.1.4), (C.1.7), and (C.1.5) equal (C.1.1), (C.1.2), and (C.1.3), respectively, and because (C.1.6) drops out ($\forall h = 1, \dots, 0$).

Induction step: Assume that the hypothesis holds for i . We next show that it will then also hold for $i + 1$. The feasibility problem of iteration $i + 1$ is obtained by applying FME on y_{1i} in (C.1.4)–(C.1.7). The constraints/rows with null coefficients for y_{1i} stay the same according to lines 9 and 10 of Algorithm 1:

$$y_{1h} \leq c_h \quad \forall h = i + 1, \dots, m \quad (\text{second constraint to last constraint of (C.1.4)}) \quad (\text{C.1.8})$$

$$-y_{1h} \leq 0 \quad \forall h = i + 1, \dots, m \quad (\text{second constraint to last constraint of (C.1.5)}) \quad (\text{C.1.9})$$

$$0 \leq c_h \quad \forall h = 1, \dots, i - 1 \quad (\text{constraints (C.1.6)})$$

The constraints with coefficient $+1$ or -1 for y_{1i} are given by

$$y_{1i} \leq c_i \text{ (first constraint of (C.1.4))} \tag{C.1.10}$$

$$-y_{1i} \leq 0 \text{ (first constraint of (C.1.5))} \tag{C.1.11}$$

$$-\sum_{h=i}^m y_{1h} \leq \sum_{h=1}^{i-1} c_h - y_1 \text{ (constraint (C.1.7)).}$$

Consequently, adding (C.1.10) and (C.1.11), as well as (C.1.10) and (C.1.7), according to lines 13 and 14 of Algorithm 1 leads to

$$0 \leq c_i \tag{C.1.12}$$

$$-\sum_{h=i+1}^m y_{1h} \leq c_i + \sum_{h=1}^{i-1} c_h - y_1 = \sum_{h=1}^i c_h - y_1. \tag{C.1.13}$$

Thus, in total, we obtain the following set of constraints from iteration i :

$$y_{1h} \leq c_h \ \forall h = i + 1, \dots, m \text{ (constraints (C.1.8))}$$

$$-y_{1h} \leq 0 \ \forall h = i + 1, \dots, m \text{ (constraints (C.1.9))}$$

$$0 \leq c_h \ \forall h = 1, \dots, i \text{ (consisting of constraints (C.1.6) and (C.1.12))}$$

$$-\sum_{h=i+1}^m y_{1h} \leq \sum_{h=1}^i c_h - y_1 \text{ (constraint (C.1.13))}$$

These constraints equal (C.1.4)–(C.1.7) with increased $i := i + 1$, which concludes the induction step.

Given this result, we subsequently consider the feasibility problem after completely executing Algorithm 1, that is, after eliminating all the distribution variables. The resulting problem is given by (C.1.6) and (C.1.7) with $i = m + 1$, because (C.1.4) and (C.1.5) drop out ($\forall h = m + 1, \dots, m$):

$$0 \leq c_h \ \forall h = 1, \dots, m \text{ (regular resources, constraints (C.1.6))}$$

$$0 \leq \sum_{h=1}^m c_h - y_1 \text{ (artificial resource, constraint (C.1.7))}$$

Therefore, we obtain a total of $\tilde{m} = 1$ artificial resource. □

C.2 Proof of Proposition 5

Proposition 5: In network type 2, the number of artificial resources is $\tilde{m} = \frac{(m-1) \cdot m}{2}$ (and thus polynomial in the number of regular resources m).

Proof: Network type 2 consists of m parallel resources and $m - 1$ flexible products. Flexible product k may be assigned to resource k or $k + 1$. Thus, the feasibility problem (5)–(7) is given by:

$$y_{11} \leq c_1 \tag{C.2.1}$$

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = 1, \dots, m - 2 \tag{C.2.2}$$

$$y_{m-1,2} \leq c_m \tag{C.2.3}$$

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = 1, \dots, m - 1 \tag{C.2.4}$$

$$-y_{k1} \leq 0 \quad \forall k = 1, \dots, m - 1 \tag{C.2.5}$$

$$-y_{k2} \leq 0 \quad \forall k = 1, \dots, m - 1 \tag{C.2.6}$$

W.l.o.g., we assume that the $2 \cdot (m - 1)$ distribution variables are eliminated in the order of $y_{11}, y_{12}, y_{21}, \dots, y_{m-1,2}$. For this purpose, we conduct the iterations $i = 1, \dots, m - 1$, with i referring to the feasibility problem after eliminating $y_{i-1,2}$ (for $i > 1$) and before eliminating y_{i1} . Please note that dummy iteration $i = 1$ corresponds to the initial feasibility problem, that each iteration comprises eliminating both distribution variables of flexible product i (and thus, two of the iterations of Algorithm 1), and that $i = m$ refers to the feasibility problem after eliminating all the distribution variables, that is, after iteration $m - 1$.

Now, by induction over i , we show the following:

Induction hypothesis: The feasibility problem in iteration $i = 1, \dots, m - 1$ is given by:

$$y_{i1} \leq c_i \quad (\text{C.2.7})$$

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = i, \dots, m-2 \quad (\text{C.2.8})$$

$$y_{m-1,2} \leq c_m \quad (\text{C.2.9})$$

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = i, \dots, m-1 \quad (\text{C.2.10})$$

$$-y_{k1} \leq 0 \quad \forall k = i, \dots, m-1 \quad (\text{C.2.11})$$

$$-y_{k2} \leq 0 \quad \forall k = i, \dots, m-1 \quad (\text{C.2.12})$$

$$0 \leq \sum_{\bar{h}=\underline{h}}^{\bar{h}} c_{\bar{h}} - \sum_{\bar{h}=\underline{h}}^{\bar{h}-1} y_{\bar{h}} \quad \forall \underline{h} = 1, \dots, i-1, \forall \bar{h} = \underline{h}, \dots, i-1 \quad (\text{C.2.13})$$

$$y_{i1} \leq \sum_{\bar{h}=\underline{h}}^i c_{\bar{h}} - \sum_{\bar{h}=\underline{h}}^{i-1} y_{\bar{h}} \quad \forall \underline{h} = 1, \dots, i-1 \quad (\text{C.2.14})$$

Induction basis: The induction hypothesis holds for $i = 1$, because (C.2.7)–(C.2.12) equal (C.2.1)–(C.2.6), respectively, and because (C.2.13) and (C.2.14) drop out ($\forall \underline{h} = 1, \dots, 0$).

Induction step: Assume that the hypothesis holds for i . We now show that it will then also hold for $i + 1$. We first eliminate y_{i1} by performing one iteration of Algorithm 1. The constraints/rows with null coefficients for y_{i1} stay the same according to lines 9 and 10 of Algorithm 1:

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = i, \dots, m-2 \quad (\text{constraints (C.2.8)})$$

$$y_{m-1,2} \leq c_m \quad (\text{constraint (C.2.9)})$$

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = i+1, \dots, m-1 \quad (\text{second constraint to last constraint of (C.2.10)}) \quad (\text{C.2.15})$$

$$-y_{k1} \leq 0 \quad \forall k = i+1, \dots, m-1 \quad (\text{second constraint to last constraint of (C.2.11)}) \quad (\text{C.2.16})$$

$$-y_{k2} \leq 0 \quad \forall k = i, \dots, m-1 \quad (\text{constraints (C.2.12)})$$

$$0 \leq \sum_{\bar{h}=\underline{h}}^{\bar{h}} c_{\bar{h}} - \sum_{\bar{h}=\underline{h}}^{\bar{h}-1} y_{\bar{h}} \quad \forall \underline{h} = 1, \dots, i-1, \forall \bar{h} = \underline{h}, \dots, i-1 \quad (\text{constraints (C.2.13)})$$

The constraints with coefficient $+1$ or -1 for y_{i1} are given by

$$y_{i1} \leq c_i \quad (\text{constraint (C.2.7)})$$

$$y_{i1} \leq \sum_{h=\underline{h}}^i c_h - \sum_{h=\underline{h}}^{i-1} y_h \quad \forall \underline{h} = 1, \dots, i-1 \text{ (constraints (C.2.14))}$$

$$-y_{i1} - y_{i2} \leq -y_i \text{ (first constraint of (C.2.10))} \quad (\text{C.2.17})$$

$$-y_{i1} \leq 0 \text{ (first constraint of (C.2.11)).} \quad (\text{C.2.18})$$

Consequently, adding (C.2.7) and (C.2.17), (C.2.7) and (C.2.18), (C.2.14) and (C.2.17), as well as (C.2.14) and (C.2.18), according to lines 13 and 14 of Algorithm 1 leads to

$$-y_{i2} \leq c_i - y_i \quad (\text{C.2.19})$$

$$0 \leq c_i \quad (\text{C.2.20})$$

$$-y_{i2} \leq \sum_{h=\underline{h}}^i c_h - \sum_{h=\underline{h}}^{i-1} y_h - y_i = \sum_{h=\underline{h}}^i c_h - \sum_{h=\underline{h}}^i y_h \quad \forall \underline{h} = 1, \dots, i-1 \quad (\text{C.2.21})$$

$$0 \leq \sum_{h=\underline{h}}^i c_h - \sum_{h=\underline{h}}^{i-1} y_h \quad \forall \underline{h} = 1, \dots, i-1. \quad (\text{C.2.22})$$

Thus, in total, we obtain the following set of constraints after the elimination of y_{i1} :

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = i, \dots, m-2 \text{ (constraints (C.2.8))}$$

$$y_{m-1,2} \leq c_m \text{ (constraint (C.2.9))}$$

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = i+1, \dots, m-1 \text{ (constraints (C.2.15))}$$

$$-y_{k1} \leq 0 \quad \forall k = i+1, \dots, m-1 \text{ (constraints (C.2.16))}$$

$$-y_{k2} \leq 0 \quad \forall k = i, \dots, m-1 \text{ (constraints (C.2.12))}$$

$$0 \leq \sum_{h=\underline{h}}^{\bar{h}} c_h - \sum_{h=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, i, \forall \bar{h} = \underline{h}, \dots, i \text{ (consisting of constraints (C.2.13), (C.2.20), and (C.2.22))} \quad (\text{C.2.23})$$

$$-y_{i2} \leq \sum_{h=\underline{h}}^i c_h - \sum_{h=\underline{h}}^i y_h \quad \forall \underline{h} = 1, \dots, i \text{ (consisting of constraints (C.2.19) and (C.2.21))} \quad (\text{C.2.24})$$

For this set of constraints, we show that, when applying another iteration of Algorithm 1 to eliminate y_{i2} , we obtain the feasibility problem of iteration $i+1$; that is (C.2.7)–(C.2.14) with increased $i := i+1$. The constraints/rows with null coefficients for y_{i2} stay the same according to lines 9 and 10 of Algorithm 1:

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = i+1, \dots, m-2 \text{ (second constraint to last constraint of (C.2.8); drops out in case that } i = m-2) \quad (\text{C.2.25})$$

$y_{m-1,2} \leq c_m$ (constraint (C.2.9); note that $i < m - 1$ according to the hypothesis; thus this constraint always has null coefficients for y_{i2})

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = i + 1, \dots, m - 1 \text{ (constraints (C.2.15))}$$

$$-y_{k1} \leq 0 \quad \forall k = i + 1, \dots, m - 1 \text{ (constraints (C.2.16))}$$

$$-y_{k2} \leq 0 \quad \forall k = i + 1, \dots, m - 1 \text{ (second constraint to last constraint of (C.2.12))} \quad (\text{C.2.26})$$

$$0 \leq \sum_{\bar{h}=\underline{h}} c_h - \sum_{\bar{h}=\underline{h}} y_h \quad \forall \underline{h} = 1, \dots, i, \forall \bar{h} = \underline{h}, \dots, i \text{ (constraints (C.2.23))}$$

The rows with coefficient +1 or -1 are

$$y_{i2} + y_{i+1,1} \leq c_{i+1} \text{ (first constraint of (C.2.8); note that } i < m - 1 \text{ according to the hypothesis, thus this constraint always exists)} \quad (\text{C.2.27})$$

$$-y_{i2} \leq 0 \text{ (first constraint of (C.2.12))} \quad (\text{C.2.28})$$

$$-y_{i2} \leq \sum_{\bar{h}=\underline{h}}^i c_h - \sum_{\bar{h}=\underline{h}}^i y_h \quad \forall \underline{h} = 1, \dots, i \text{ (constraints (C.2.24)).}$$

Consequently, adding (C.2.27) and (C.2.28), as well as (C.2.27) and (C.2.24), according to lines 13 and 14 of Algorithm 1 leads to

$$y_{i+1,1} \leq c_{i+1} \quad (\text{C.2.29})$$

$$y_{i+1,1} \leq c_{i+1} + \sum_{\bar{h}=\underline{h}}^i c_h - \sum_{\bar{h}=\underline{h}}^i y_h = \sum_{\bar{h}=\underline{h}}^{i+1} c_h - \sum_{\bar{h}=\underline{h}}^i y_h \quad \forall \underline{h} = 1, \dots, i. \quad (\text{C.2.30})$$

Thus, in total, we obtain the following set of constraints after the elimination of y_{i2} :

$$y_{i+1,1} \leq c_{i+1} \text{ (constraint (C.2.29))}$$

$$y_{k2} + y_{k+1,1} \leq c_{k+1} \quad \forall k = i + 1, \dots, m - 2 \text{ (constraints (C.2.25))}$$

$$y_{m-1,2} \leq c_m \text{ (constraint (C.2.9))}$$

$$-y_{k1} - y_{k2} \leq -y_k \quad \forall k = i + 1, \dots, m - 1 \text{ (constraints (C.2.15))}$$

$$-y_{k1} \leq 0 \quad \forall k = i + 1, \dots, m - 1 \text{ (constraints (C.2.16))}$$

$$-y_{k2} \leq 0 \quad \forall k = i + 1, \dots, m - 1 \text{ (constraints (C.2.26))}$$

$$0 \leq \sum_{\bar{h}=\underline{h}} \bar{c}_h - \sum_{\bar{h}=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, i, \forall \bar{h} = \underline{h}, \dots, i \text{ (constraints (C.2.23))}$$

$$y_{i+1,1} \leq \sum_{\bar{h}=\underline{h}}^{i+1} c_h - \sum_{\bar{h}=\underline{h}}^i y_h \quad \forall \underline{h} = 1, \dots, i \text{ (constraints (C.2.30))}$$

These constraints equal (C.2.7)–(C.2.14) with increased $i := i + 1$, which concludes the induction step. Given this result, we can now formally state the feasibility problem after performing $m - 2$ iterations by simply setting $i = m - 1$ in the hypothesis. We obtain:

$$y_{m-1,1} \leq c_{m-1}$$

$$y_{m-1,2} \leq c_m$$

$$-y_{m-1,1} - y_{m-1,2} \leq -y_{m-1}$$

$$-y_{m-1,1} \leq 0$$

$$-y_{m-1,2} \leq 0$$

$$0 \leq \sum_{\bar{h}=\underline{h}} \bar{c}_h - \sum_{\bar{h}=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, m - 2, \forall \bar{h} = \underline{h}, \dots, m - 2$$

$$y_{m-1,1} \leq \sum_{\bar{h}=\underline{h}}^{m-1} c_h - \sum_{\bar{h}=\underline{h}}^{m-2} y_h \quad \forall \underline{h} = 1, \dots, m - 2$$

Finally, we perform the remaining $m - 1$ -th iteration, that is, two final iterations of Algorithm 1 to subsequently eliminate $y_{m-1,1}$ and $y_{m-1,2}$. After the elimination of $y_{m-1,1}$, we obtain the following set of constraints:

$$y_{m-1,2} \leq c_m$$

$$-y_{m-1,2} \leq 0$$

$$0 \leq \sum_{h=\underline{h}}^{\bar{h}} c_h - \sum_{h=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, m-1 \quad \forall \bar{h} = \underline{h}, \dots, m-1$$

$$-y_{m-1,2} \leq \sum_{h=\underline{h}}^{m-1} c_h - \sum_{h=\underline{h}}^{m-1} y_h \quad \forall \underline{h} = 1, \dots, m-1$$

After eliminating $y_{m-1,2}$ —that is, after eliminating all the distribution variables of the original feasibility problem—we obtain the constraints

$$0 \leq \sum_{h=\underline{h}}^{\bar{h}} c_h - \sum_{h=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, m \quad \forall \bar{h} = \underline{h}, \dots, m,$$

which may be rewritten as:

$$0 \leq c_h \quad \forall h = 1, \dots, m \text{ (regular resources)}$$

$$0 \leq \sum_{h=\underline{h}}^{\bar{h}} c_h - \sum_{h=\underline{h}}^{\bar{h}-1} y_h \quad \forall \underline{h} = 1, \dots, m-1, \forall \bar{h} = \underline{h} + 1, \dots, m \text{ (artificial resources)}$$

Therefore, we have a total of $\tilde{m} = \sum_{h=1}^{m-1} h = \frac{(m-1)m}{2}$ artificial resources. \square

D Technical details of the implemented methods

In Section 5, we evaluated the average revenues obtained by using our surrogate approach (*DPD-surr*) in comparison with two revenue management methods (*DPD-ah* and *CDLP-surr*), as well as an upper bound (*UB*) on the optimal expected revenue of DP-flex (1). In the following, we provide the technical details.

D.1 Upper bound (UB)

As the upper bound on the optimal expected revenue of (1), we use the optimal objective value of the corresponding CDLP formulation, which Gallego et al. (2004) propose (**CDLP-flex**):

$$\text{Maximize } \sum_S t(S) \cdot \lambda \cdot \left(\sum_j P_j^{reg}(S) \cdot r_j^{reg} + \sum_k P_k^{flex}(S) \cdot r_k^{flex} \right) \quad (\text{D.1.1})$$

subject to

$$\sum_S t(S) \cdot \lambda \cdot \sum_j P_j^{reg}(S) \cdot a_{hj} + \sum_k \sum_{j \in \mathcal{M}_k} a_{hj} \cdot y_{kj} \leq c_h \quad \forall h \quad (\text{D.1.2})$$

$$\sum_S t(S) \cdot \lambda \cdot P_k^{flex} = \sum_{j \in \mathcal{M}_k} y_{kj} \quad \forall k \quad (\text{D.1.3})$$

$$\sum_S t(S) = T \quad (\text{D.1.4})$$

$$t(S) \geq 0 \quad \forall S \quad (\text{D.1.5})$$

$$y_{kj} \geq 0 \forall k, j \in \mathcal{M}_k \quad (\text{D.1.6})$$

In this model, the variable $t(S)$ denotes how long set S is offered. Please note that constraints (D.1.2), (D.1.3), and (D.1.6), after some minor rearrangements, equal the (relaxed) feasibility problem (5)–(7).

In order to solve CDLP-flex and obtain UB , we use column generation (see, e.g., Liu and van Ryzin (2008) and Miranda Bront et al. (2009) for an extensive description in the context of standard revenue management). We start with a reduced number of columns in CDLP-flex; that is, with only a subset of the possible offer sets. Let π_h^{red} , μ_k^{red} , and σ^{red} denote the optimal dual prices for restrictions (D.1.2), (D.1.3), and (D.1.4) of this reduced problem. Thereafter, we have to check whether there is an offer set with positive reduced costs that must be included. More precisely, a column corresponding to a new offer set is the optimal solution of the following column generation sub-problem:

$$\begin{aligned} & \max_S \{ \lambda \cdot (\sum_j P_j^{reg}(S) \cdot (r_j^{reg} - \sum_h a_{hj} \cdot \pi_h^{red}) \\ & + \sum_k P_k^{flex}(S) \cdot (r_k^{flex} - \mu_k^{red})) \} - \sigma^{red} \end{aligned} \quad (\text{D.1.7})$$

The second term in the argument of the maximum function is due to the consideration of flexible products. The solution technique depends on the choice model used.

Please note that we can obtain UB alternatively by using the surrogate reformulation from Section 4. To see this, we apply the surrogate network on the standard CDLP formulation without flexible products (see, e.g., Liu and van Ryzin (2008)) and obtain the following formulation (**CDLP-surr**):

$$\text{Maximize } \sum_S t(S) \cdot \lambda \cdot (\sum_j P_j^{reg}(S) \cdot r_j^{reg} + \sum_{k \in S} P_k^{flex}(S) \cdot r_k^{flex}) \quad (\text{D.1.8})$$

subject to

$$\sum_S t(S) \cdot \lambda \cdot \sum_j P_j^{reg}(S) \cdot a_{hj} \leq c_h \quad \forall h \quad (\text{D.1.9})$$

$$\sum_S t(S) \cdot \lambda \cdot (\sum_j P_j^{reg}(S) \cdot \tilde{a}_{ij} + \sum_k P_k^{flex}(S) \cdot \tilde{b}_{ik}) \leq \tilde{c}_i \quad \forall i \quad (\text{D.1.10})$$

$$\sum_S t(S) = T \quad (\text{D.1.11})$$

$$t(S) \geq 0 \quad \forall S \quad (\text{D.1.12})$$

A comparison of CDLP-flex with CDLP-surr shows that they only differ in the constraints representing the feasibility problem: While CDLP-flex contains the original feasibility problem, CDLP-surr contains the transformed feasibility problem (constraints (D.1.9) and (D.1.10)). Thus, both CDLPs are equivalent. This result is intuitive, because both the original and the surrogate networks represent the same stochastic problem (represented by the DPs). In this sense, this result for the deterministic equivalent (given by the CDLP) of the stochastic problem parallels the result obtained in Section 4.2 in respect of the DPs. Please note that Condition 1 is not required for equivalence here.

D.2 Ad hoc assignment DPD approach (DPD-ah)

This method is based on an idea that was already investigated by Steinhardt and Gönsch (2012) in respect of the special case of upgrades and without customer choice. It forgoes flexibility and immediately assigns flexible products (ad hoc) after sale. As this assignment is irrevocable, we can immediately reduce the remaining capacity and do not need to store any commitments. Thus, a resource-based state space is obtained, and DPD by resources is possible. However, existing choice-based approaches do not include the described ad hoc assignment of flexible products and have to be modified appropriately. In the following, we carry out these modifications on the DPD approach of Liu and van Ryzin (2008).

The immediate assignment of a flexible product k to the current best of its alternatives $j \in \mathcal{M}_k$ is captured by the second line in the Bellman equation

$$\begin{aligned} V_t^{ah}(\mathbf{c}) = & \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} + V_{t-1}^{ah}(\mathbf{c} - \mathbf{a}_j) \right) \right. \\ & + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \max_{j \in \mathcal{M}_k} \left\{ r_k^{flex} + V_{t-1}^{ah}(\mathbf{c} - \mathbf{a}_j) \right\} \\ & \left. + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}^{ah}(\mathbf{c}) \right\} \end{aligned} \quad (\text{D.2.1})$$

with the boundary conditions $V_t^{ah}(\mathbf{c}) = -\infty$ if $\mathbf{c} \not\geq \mathbf{0}$ and $V_0^{ah}(\mathbf{c}) = 0 \forall \mathbf{c} \geq \mathbf{0}$.

The standard starting point of the decomposition is the corresponding CDLP formulation CDLP-flex (D.1.1)–(D.1.6). Let π_h^{ah} denote the optimal dual prices associated with

the capacity of resource h (constraint (D.1.2)). We then obtain the following one-dimensional problem to assess the value of capacity of each resource $h' \in \mathcal{H}$:

$$\begin{aligned} V_t^{ah,h'}(c_{h'}) &= \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} - \sum_{h \neq h'} a_{hj} \cdot \pi_h^{ah} + V_{t-1}^{ah,h'}(c_{h'} - a_{h'j}) \right) \right. \\ &+ \sum_k \lambda \cdot P_k^{flex}(S) \cdot \max_{j \in \mathcal{M}_k} \left\{ r_k^{flex} - \sum_{h \neq h'} a_{hj} \cdot \pi_h^{ah} + V_{t-1}^{ah,h'}(c_{h'} - a_{h'j}) \right\} \\ &\left. + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}^{ah,h'}(c_{h'}) \right\} \end{aligned} \quad (D.2.2)$$

with boundary conditions $V_t^{ah,h'}(c_{h'}) = -\infty$ if $c_{h'} < 0$ and $V_0^{ah,h'}(c_{h'}) = 0 \forall c_{h'} \geq 0$.

During the booking horizon (that is, during the simulations), we approximate the opportunity cost of a regular product j and all flexible products' alternatives $j \in \mathcal{M}_k$, respectively, as the sum of the required resources' opportunity cost. More formally, with resource h 's opportunity cost defined as

$$\Delta_h V_t^{ah,h}(c_h) := V_t^{ah,h}(c_h) - V_t^{ah,h}(c_h - 1), \quad (D.2.3)$$

the offer set is the optimal solution of

$$\begin{aligned} \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot \left(r_j^{reg} - \sum_h a_{hj} \cdot \Delta_h V_{t-1}^{ah,h}(c_h) \right) \right. \\ \left. + \sum_k \lambda \cdot P_k^{flex}(S) \cdot \max_{j \in \mathcal{M}_k} \left\{ r_k^{flex} - \sum_h a_{hj} \cdot \Delta_h V_{t-1}^{ah,h}(c_h) \right\} \right\}. \end{aligned} \quad (D.2.4)$$

Note that compared to the standard setting described, for example, in Liu and van Ryzin (2008), the second lines in equations (D.2.2) and (D.2.4) are extensions that are due to the consideration of flexible products. These modifications follow the lines of the modifications performed without customer choice in Steinhardt and Gönsch (2012) in respect of upgrades, and in Gönsch and Steinhardt (2013) in respect of opaque products. Similar to the column generation sub-problems used to solve CDLP-flex and CDLP-surr, the solution technique applied to find the optimal offer set S in (D.2.2) and (D.2.4) depends on the choice model used.

D.3 Primal solution of CDLP-surr (CDLP-surr)

This method operationalizes the optimal primal solution of CDLP-surr (D.1.8)–(D.1.12). Recall that the optimal solution gives us the time a set S is offered. Alternatively, the same solution is obtained by CDLP-flex (D.1.1)–(D.1.6).

We round fractional values of the decision variables $t(S)$ to the nearest integer. The sequence in which we offer the sets follows the lexicographic order in which the sets appear in the optimal solution. Please note that the offer sets are static over a number of periods. Thus, we have to check continuously (i.e., in each period) whether the capacity allows for offering the products contained in the static set. To do so, we use the capacity check of the surrogate network (i.e., constraints (11) and (12)). If the capacity is not sufficient to sell a product, it is removed from the set. Note that it is also possible to check capacity by solving the original feasibility problem (2)–(4).

D.4 Surrogate DPD approach (DPD-surr)

Our main method *DPD-surr* is obtained by applying the surrogate network to the DPD approach of Liu and van Ryzin (2008).

Analogously to *DPD-ah*, the starting point of the decomposition is the corresponding CDLP formulation; that is, CDLP-surr (D.1.8)–(D.1.12). Let π_h denote the optimal dual prices of regular resource h (constraint (D.1.9)). Furthermore, let $\tilde{\pi}_i$ denote the optimal dual prices of artificial resource i (constraint (D.1.10)).

We subsequently obtain the following two types of one-dimensional problems to assess the value of capacity of resources $h' \in \mathcal{H}$ and $i' \in \tilde{\mathcal{H}}$:

$$V_t^{surr,h'}(c_{h'}) = \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot (r_j^{reg} - \sum_{h \neq h'} a_{hj} \cdot \pi_h - \sum_i \tilde{a}_{ij} \cdot \tilde{\pi}_i + V_{t-1}^{surr,h'}(c_{h'} - a_{h'j})) + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot V_{t-1}^{surr,h'}(c_{h'}) \right\} \quad (D.4.1)$$

with boundary conditions $V_t^{surr,h'}(c_{h'}) = -\infty$ if $c_{h'} < 0$ and $V_0^{surr,h'}(c_{h'}) = 0 \forall c_{h'} \geq 0$ and

$$\begin{aligned} \tilde{V}_t^{surr,i'}(\tilde{c}_{i'}) = \max_S \left\{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot (r_j^{reg} - \sum_h a_{hj} \cdot \pi_h - \sum_{i \neq i'} \tilde{a}_{ij} \cdot \tilde{\pi}_i + \tilde{V}_{t-1}^{surr,i'}(\tilde{c}_{i'} - \tilde{a}_{i'j})) \right. \\ \left. + \sum_k \lambda \cdot P_k^{flex}(S) \cdot (r_k^{flex} - \sum_{i \neq i'} \tilde{b}_{ik} \cdot \tilde{\pi}_i + \tilde{V}_{t-1}^{surr,i'}(\tilde{c}_{i'} - \tilde{b}_{i'k})) + (\lambda \cdot P_0(S) + 1 - \lambda) \cdot \tilde{V}_{t-1}^{surr,i'}(\tilde{c}_{i'}) \right\} \end{aligned} \quad (D.4.2)$$

with boundary conditions $\tilde{V}_t^{surr,i'}(\tilde{c}_{i'}) = -\infty$ if $\tilde{c}_{i'} < 0$ and $\tilde{V}_0^{surr,i'}(\tilde{c}_{i'}) = 0 \forall \tilde{c}_{i'} \geq 0$.

During the booking horizon, we approximate the opportunity cost as the sum of the required resources' opportunity cost. More formally, let regular resource h 's opportunity cost be defined as

$$\Delta_h V_t^{surr,h}(c_h) := V_t^{surr,h}(c_h) - V_t^{surr,h}(c_h - 1) \quad (\text{D.4.3})$$

and artificial resource i 's opportunity cost be defined as

$$\tilde{\Delta}_i V_t^{surr,i}(c_h) := \tilde{V}_t^{surr,i}(c_i) - \tilde{V}_t^{surr,i}(c_i - 1). \quad (\text{D.4.4})$$

Then the offer set is the optimal solution of

$$\begin{aligned} & \max_S \{ \sum_j \lambda \cdot P_j^{reg}(S) \cdot (r_j^{reg} - \sum_h a_{hj} \cdot \Delta_h V_{t-1}^{surr,h}(c_h) - \sum_i \tilde{a}_{ij} \cdot \tilde{\Delta}_i V_{t-1}^{surr,i}(c_h)) \\ & + \sum_k \lambda \cdot P_k^{flex}(S) \cdot (r_k^{flex} - \sum_i \tilde{b}_{ik} \cdot \tilde{\Delta}_i V_{t-1}^{surr,i}(c_h)) \}. \end{aligned} \quad (\text{D.4.5})$$

E Detailed values for products and segments in network 2

Regarding network 2, Table E.1 and Table E.2 summarize regular products' capacity consumption and revenues, as well as the segments' arrival probabilities, consideration sets, and preference weights.

Product	Legs	Revenue	Product	Legs	Revenue
1	1	1000	12	1	500
2	2	400	13	2	200
3	3	400	14	3	200
4	4	300	15	4	150
5	5	300	16	5	150
6	6	500	17	6	250
7	7	500	18	7	250
8	(2,4)	600	19	(2,4)	300
9	(3,5)	600	20	(3,5)	300
10	(2,6)	700	21	(2,6)	350
11	(3,7)	700	22	(3,7)	350

Table E.1: Description of regular products for network 2

Segment	Class	Probability	Consideration set	Preference vector
1	H	0.08	{1,8,9}	(10,5,5)
2	L	0.16	{12,19,20}	(10,10,5)
3	H	0.05	{2,3}	(10,10)
4	L	0.16	{13,14}	(10,10)
5	H	0.10	{4,5}	(10,10)
6	L	0.12	{15,16}	(10,5)
7	H	0.02	{6,7}	(10,5)
8	L	0.04	{17,18}	(10,10)
9	H	0.02	{10,11}	(10,5)
10	L	0.04	{21,22}	(10,10)
11	-	0.05	{f1}	(10)
12	-	0.02	{f2}	(10)
13	-	0.05	{f3}	(10)
14	-	0.04	{f4}	(10)
15	-	0.02	{f5}	(10)

Table E.2: Descriptions of customer segments for network 2

Acknowledgements

We thank the Associate Editor, as well as two anonymous reviewers, for the valuable feedback that helped improve the paper.

References

- Adelman D (2007) Dynamic bid prices in revenue management. *Operations Research* 55(4): 647–661.
- Bartodziej P, Derigs U, Zils M (2006) O&D revenue management in cargo airlines - A mathematical programming approach. *OR Spectrum* 29(1): 105–121.
- Belobaba PP (1987) Air travel demand and airline seat inventory management. Dissertation, Massachusetts Institute of Technology, Cambridge.
- Belobaba PP (1989) Application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37(2): 183–197.
- Bertsimas D, Popescu I (2003) Revenue management in a dynamic network environment. *Transportation Science* 37(3): 257–277.
- Bertsimas D, Tsitsiklis JN (1997) Introduction to linear optimization. 3rd ed. Athena Scientific, Belmont, Massachusetts.
- Birbil Şİ, Frenk JBG, Gromicho JAS, Zhang S (2014) A network airline revenue management framework based on decomposition by origins and destinations. *Transportation Science* 48(3): 313–333.
- Chen D (1998) Network flows in hotel yield management. Working paper, Cornell University, New York.
- Chen S, Gallego G, Li MZ, Lin B (2010) Optimal seat allocation for two-flight problems with a flexible demand segment. *European Journal of Operational Research* 201(3): 897–908.
- Cooper WL, Homem-de-Mello T (2007) Some decomposition methods for revenue management. *Transportation Science* 41(3): 332–353.
- Davis JM, Gallego G, Topaloglu H (2014) Assortment optimization under variants of the nested logit model. *Operations Research* 62(2): 250–273.
- Erdelyi A, Topaloglu H (2009) Separable approximations for joint capacity control and overbooking decisions in network revenue management. *Journal of Revenue and Pricing Management* 8(1): 3–20.
- Erdelyi A, Topaloglu H (2010) A dynamic programming decomposition method for making overbooking decisions over an airline network. *INFORMS Journal on Computing* 22(3): 443–456.

- Fink A, Reiners T (2006) Modeling and solving the short-term car rental logistics problem. *Transportation Research Part E: Logistics and Transportation Review* 42(4): 272–292.
- Gallego G, Iyengar G, Phillips RL, Dubey A (2004) Managing flexible products on a network. Working paper, Columbia University, New York.
- Gallego G, Phillips RL (2004) Revenue management of flexible products. *Manufacturing & Service Operations Management* 6(4): 321–337.
- Gallego G, Ratliff R, Shebalov S (2015) A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research* 63(1): 212–232.
- Gallego G, Stefanescu C (2009) Upgrades, upsells and pricing in revenue management. Working paper, Columbia University, New York.
- Geraghty MK, Johnson E (1997) Revenue management saves national car rental. *Interfaces* 27(1): 107–127.
- Glover F, Glover R, Lorenzo J, McMillan C (1982) The passenger-mix problem in the scheduled airlines. *Interfaces* 12(3): 73–80.
- Gönsch J, Steinhardt C (2013) Using dynamic programming decomposition for revenue management with opaque products. *BuR - Business Research* 6(1): 94–115.
- Gönsch J, Steinhardt C (2015) On the incorporation of upgrades into airline network revenue management. *Review of Managerial Science* 9(4): 635–660.
- Hosseinalifam M (2014) A mathematical programming framework for network capacity control in customer choice-based revenue management. Dissertation, École polytechnique de Montréal, Montréal.
- Jordan WC, Graves SC (1995) Principles on the benefits of manufacturing process flexibility. *Management Science* 41(4): 577–594.
- Kemmer P, Strauss AK, Winter T (2011) Dynamic simultaneous fare proration for large-scale network revenue management. *Journal of the Operational Research Society* 63(10): 1336–1350.
- Kimms A, Müller-Bungart M (2007) Revenue management for broadcasting commercials: the channel’s problem of selecting and scheduling the advertisements to be aired. *International Journal of Revenue Management* 1(1): 28–44.
- Kunnumkal S, Topaloglu H (2008) A tractable revenue management model for capacity allocation and overbooking over an airline network. *Flexible Services and Manufacturing Journal* 20(3): 125–147.
- Kunnumkal S, Topaloglu H (2010) A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior. *Production and Operations Management* 19(5): 575–590.

- Lautenbacher CJ, Stidham S (1999) The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science* 33(2): 136–146.
- Leachman RC, Carmon TF (1992) On capacity modeling for production planning with alternative machine types. *IIE Transactions* 24(4): 62–72.
- Lee TC, Hersh M (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27(3): 252–265.
- Littlewood K (1972) Forecasting and control of passenger bookings. *Proceedings of the 12th AGIFORS Symposium, Netanja*: 95–117.
- Liu Q, van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2): 288–310.
- Martin RK (1999) *Large scale linear and integer optimization: A unified approach*. Kluwer, Boston.
- Meissner J, Strauss A (2012) Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research* 216(2): 459–468.
- Meissner J, Strauss A, Talluri KT (2013) An enhanced concave program relaxation for choice network revenue management. *Production and Operations Management* 22(1): 71–87.
- Miranda Bront JJ, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Operations Research* 57(3): 769–784.
- Oosten M (2004) Revenue management with flexible due dates. Presentation, Conference of POMS College of Service Operations, New York.
- Pachon JE, Iakovou E, Ip C, Aboudi R (2003) A synthesis of tactical fleet planning models for the car rental industry. *IIE Transactions* 35(9): 907–916.
- Paulraj S, Sumathi P (2010) A comparative study of redundant constraints identification methods in linear programming problems. *Mathematical Problems in Engineering* 2010(1): 1–16.
- Petrick A, Gönsch J, Steinhardt C, Klein R (2010) Dynamic control mechanisms for revenue management with flexible products. *Computers & Operations Research* 37(11): 2027–2039.
- Petrick A, Steinhardt C, Gönsch J, Klein R (2012) Using flexible products to cope with demand uncertainty in revenue management. *OR Spectrum* 34(1): 215–242.
- Schrijver A (1998) *Theory of linear and integer programming*. Wiley, Chichester.
- Shumsky RA, Zhang F (2009) Dynamic capacity management with substitution. *Operations Research* 57(3): 671–684.

- Steinhardt C, Gönsch J (2012) Integrated revenue management approaches for capacity control with planned upgrades. *European Journal of Operational Research* 223(2): 380–391.
- Strauss AK, Talluri KT (2015) Tractable consideration set structures for network revenue management. Working paper, University of Warwick, Warwick.
- Talluri KT (2001) Airline revenue management with passenger routing control: A new model with solution approaches. *International Journal of Services Technology and Management* 2(1/2): 102–115.
- Talluri KT, van Ryzin G (1998) An analysis of bid-price controls for network revenue management. *Management Science* 44(11): 1577–1593.
- Talluri KT, van Ryzin G (2004a) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.
- Talluri KT, van Ryzin G (2004b) *The theory and practice of revenue management*. Springer, New York.
- Tong C, Topaloglu H (2014) On the approximate linear programming approach for network revenue management problems. *INFORMS Journal on Computing* 26(1): 121–134.
- Train K (2009) *Discrete choice methods with simulation*. 2nd ed. Cambridge University Press, Cambridge.
- Vossen TWM, Zhang D (2015a) A dynamic disaggregation approach to approximate linear programs for network revenue management. *Production and Operations Management* 24(3): 469–487.
- Vossen TWM, Zhang D (2015b) Reductions of approximate linear programs for network revenue management. *Operations Research* 63(6): 1352–1371.
- Walter M, Truemper K (2013) Implementation of a unimodularity test. *Mathematical Programming Computation* 5(1): 57–73.
- Williams HP, Hooker JN (2016) Integer programming as projection. Appears in: *Discrete Optimization*.
- Zhang D (2011) An improved dynamic programming decomposition approach for network revenue management. *Manufacturing & Service Operations Management* 13(1): 35–52.
- Zhang D, Adelman D (2009) An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Science* 43(3): 381–394.

Beitrag B3:
Revenue management with flexible products:
The value of flexibility and its incorporation into DLP-
based approaches

Jochen Gönsch ^a, Sebastian Koch ^b, Claudius Steinhardt ^c

(^a Universität Duisburg-Essen, ^b Universität Augsburg, ^c Universität der Bundeswehr München)

International Journal of Production Economics 153(2014): 280–294.

DOI: <https://doi.org/10.1016/j.ijpe.2014.03.010>

Beitrag B4:
**An EMSR-based approach for revenue management
with integrated upgrade decisions**

Jochen Gönsch ^a, Sebastian Koch ^b, Claudius Steinhardt ^c

(^a Universität Duisburg-Essen, ^b Universität Augsburg, ^c Universität der Bundeswehr München)

Computers & Operations Research 40(10): 2532–2542.

DOI: <https://doi.org/10.1016/j.cor.2013.01.012>

Beitrag B5:

Practical decision rules for risk-averse revenue management using simulation-based optimization

Sebastian Koch ^a, Jochen Gönsch ^b, Michael Hassler ^a, Robert Klein ^a

(^a Universität Augsburg, ^b Universität Duisburg-Essen)

Journal of Revenue and Pricing Management 15(6): 468–487.

DOI: <https://doi.org/10.1057/s41272-016-0065-x>

Abstract

In practice, human-decision makers often feel uncomfortable with the risk-neutral revenue management systems' output. Reasons include a low number of repetitions of similar events, a critical impact of the achieved revenue for economic survival, or simply business constraints imposed by management. However, solving capacity control problems is a challenging task for many risk measures and the approaches are often not compatible with existing software systems.

In this paper, we propose a flexible framework for risk-averse capacity control under customer choice behavior. Existing risk-neutral decision rules are augmented by the integration of adjustable parameters. Our key idea is the application of simulation-based optimization (SBO) to calibrate these parameters. This allows to easily tailor the resulting capacity control mechanism to almost every risk measure and customer choice behavior.

In an extensive simulation study, we analyze the impact of our approach on expected utility, conditional value-at-risk (CVaR), and expected value. The results show a superior performance in comparison to risk-neutral approaches from literature.

Keywords: Revenue Management, Capacity Control, Risk-Aversion, Conditional Value-at-Risk

1 Introduction

During the last decades, revenue management has become one of the most successful fields of application for operations research in practice. Its main task is capacity control, which is usually described as controlling the availability of differentiated products over a given booking horizon such that the expected revenue is maximized. The assumption of risk-neutrality lies at the heart of this classical definition and is justified by a large number of repetitions of similar decision problems. However, human decision makers, who tend to be risk-averse, often doubt this assumption. In daily practice, they feel uncomfortable with the capacity control system's output and overwrite it manually with less aggressive decisions. Furthermore, in many fields of application, the number of repetitions is too small to justify the use of expected value and a single event is critical for economic survival. Risk-aversion first became popular in economics and finance, but it is today also increasingly considered in revenue management. The underlying trade-off is to give up a portion of expected value in order to reduce the risk of poor outcomes.

The problem of risk-averse capacity control can be solved to optimality by dynamic programming (DP). However, building a DP formulation is a challenging task for many risk measures. In many cases, the state space must be augmented and the resulting DP formulation becomes intractable. Furthermore, DP formulations are not compatible with many existing revenue management systems.

Our main contribution is to propose a flexible framework for risk-averse capacity control. In practice, revenue management systems are fixed in the long run and the capacity control process is modeled by standard decision rules such as bid prices. Therefore, our framework is based on the risk-neutral formulation. Risk-aversion is then integrated by augmenting existing capacity control mechanisms with a few parameters that can be calibrated. Existing research recommends that this is done manually by human decision makers. However, we suggest the use of simulation-based optimization (SBO) which allows automated optimization and a higher number of parameters. The resulting approach is quite general. It can be used with arbitrary demand models, risk measures, and

network structures. In an extensive simulation study, we illustrate the impact of our approach on expected utility, conditional value-at-risk (CVaR), and expected revenue in various settings with customer choice and different network structures.

The remainder of this paper is structured as follows: First, we restate the risk-neutral problem of capacity control under customer choice behavior, review the relevant scientific literature and position our work. Based upon this, we present our framework for risk-averse capacity control, including a detailed description of the components. We continue with the simulation study, followed by a discussion of the results and a conclusion.

2 Background and previous research

Research from three areas of revenue management is relevant for our work. First, we restate the problem of risk-neutral capacity control under customer choice behavior and summarize standard solution approaches. Then, we discuss research on risk-averse capacity control and the use of SBO for capacity control.

2.1 Risk-neutral capacity control under customer choice-behavior

Initially, revenue management (RM) was based on the well-known independent demand assumption. Overviews can be found in the textbooks of Talluri and van Ryzin (2004b) and Phillips (2005).

Later, research considered that most customers actually choose between several more or less suitable products. Gallego et al. (2004), Talluri and van Ryzin (2004a), and Liu and van Ryzin (2008) established capacity control under a general discrete choice model of demand. In this setting, a firm disposes of resources $i = 1, \dots, m$ which are jointly used by products $j = 1, \dots, n$. The products are associated with revenues $\mathbf{r} = (r_1, \dots, r_n)^T$. Furthermore, each product j has a capacity consumption $\mathbf{a}_j = (a_{1j}, \dots, a_{mj})^T$, which is either $a_{ij} = 1$ if product j requires resource i or $a_{ij} = 0$ else. Resources' remaining capacity is denoted by the vector $\mathbf{c} = (c_1, \dots, c_m)^T$, the initial endowment is given by

$\mathbf{c}^0 = (c_1^0, \dots, c_m^0)^T$. Customers arrive successively and stochastically over time. The booking horizon is discretized into sufficiently small time periods $t = 1, \dots, T$, such that in each period t at most one customer arrives. Thus, at most one product can be sold in each period. The periods are numbered forward in time. Any capacity remaining at the end of the booking horizon is worthless and overbooking of the given resources' capacity is not allowed.

In each period t , the firm's risk-neutral decision problem is to determine a subset of products to offer, called the offer set, so that the overall expected revenue $V_1(\mathbf{c}^0)$ is maximized. The offer set is captured by the vector $\mathbf{x} = (x_1, \dots, x_n)^T$ of binary decision variables with $x_j = 1$ if product j is offered for sale. Product j is sold with probability $p_{tj}(\mathbf{x})$ and no purchase is made with probability $p_{t0}(\mathbf{x})$.

Let the value function $V_t(\mathbf{c})$ denote the optimal expected revenue-to-go in period t with capacity \mathbf{c} and let $\Delta_j V_t(\mathbf{c}) := V_t(\mathbf{c}) - V_t(\mathbf{c} - \mathbf{a}_j)$ denote the opportunity cost of selling one unit of product j . Then, $V_t(\mathbf{c})$ and the expected revenue-maximizing offer set can be computed recursively by the following DP formulation (DP-EV)

$$V_t(\mathbf{c}) = \max_{\mathbf{x}} \left\{ \sum_{j=1}^n p_{tj}(\mathbf{x}) \cdot (r_j - \Delta_j V_{t+1}(\mathbf{c})) \right\} + V_{t+1}(\mathbf{c}) \quad (1)$$

subject to the boundary conditions $V_t(\mathbf{c}) = -\infty$ if $\mathbf{c} \not\geq \mathbf{0}$ and $V_{T+1}(\mathbf{c}) = 0$ if $\mathbf{c} \geq \mathbf{0}$.

Two issues render DP-EV difficult to solve optimally: recursively calculating the opportunity cost $\Delta_j V_t(\mathbf{c})$ and solving the maximization over all 2^n possible offer sets. Over time, different heuristic approaches have been developed. Regarding the first issue, virtually all approaches use additive bid prices π_{tic_i} that reflect the current value of one unit of capacity of resource i in period t with remaining capacity c_i . With these values, an approximation $\tilde{\Delta}_j V_t(\mathbf{c})$ of the opportunity cost can be obtained:

$$\tilde{\Delta}_j V_t(\mathbf{c}) = \sum_{i=1}^m a_{ij} \cdot \pi_{tic_i} \quad (2)$$

The approaches differ in how the bid prices are computed, but the main idea is usually to derive an easy-to-compute upper bound on $V_t(\mathbf{c})$ and use information from this upper bound to approximate the opportunity cost in an offline stage (that is, before the booking horizon starts). Such approximations can be found, for example, in Liu and van

Ryzin (2008), Miranda Bront et al. (2009), Zhang and Adelman (2009), and Meissner and Strauss (2012b). Online (that is, during the booking horizon), the offer set is then determined by solving the maximization, which is an assortment optimization problem (see, e.g., Miranda Bront et al. (2009)):

$$\max_{\mathbf{x}} \left\{ \sum_{j=1}^n p_{tj}(\mathbf{x}) \cdot (r_j - \tilde{\Delta}_j V_{t+1}(\mathbf{c})) \right\} \quad (3)$$

The technique used to solve (3) strongly depends on the choice model assumed. For example, under the independent demand model, (3) reduces to the classical method of simply offering all products for which revenue exceeds (an approximation of) opportunity cost:

$$r_j \geq \tilde{\Delta}_j V_{t+1}(\mathbf{c}) \quad (4)$$

A popular way to manage the selling process, in particular in practice, is to use this independent demand decision rule (4) in combination with additive bid prices (2). This kind of capacity control approach is often referred to as bid price control. Even if demand is not independent, decision rule (4) can be used heuristically (see, e.g., Chaneton and Vulcano (2011) and Meissner and Strauss (2012a)).

Table 1 summarizes the notation used throughout this section.

$i = 1, \dots, m$	resources	$p_{tj}(\mathbf{x})$	purchase probability of product j given offer set \mathbf{x}
$j = 1, \dots, n$	products	$p_{t0}(\mathbf{x})$	no-purchase probability given offer set \mathbf{x}
$t = 1, \dots, T$	time periods (numbered forward)	$V_t(\mathbf{c})$	optimal expected revenue-to-go in period t with capacity \mathbf{c}
$\mathbf{c} = (c_1, \dots, c_m)^T$	remaining capacity	$\Delta_j V_t(\mathbf{c})$	opportunity cost of product j
$\mathbf{c}^0 = (c_1^0, \dots, c_m^0)^T$	initial capacity	$\tilde{\Delta}_j V_t(\mathbf{c})$	approximation of opportunity cost of product j
$\mathbf{r} = (r_1, \dots, r_n)^T$	product revenues	π_{tic_i}	bid price of resource i in period t with remaining capacity c_i
$\mathbf{a}_j = (a_{1j}, \dots, a_{mj})^T$	capacity consumption of product j		
$\mathbf{x} = (x_1, \dots, x_n)^T$	offer set of products		

Table 1: Notation introduced in this section

2.2 Risk-averse capacity control

In this section, we briefly outline the consideration of risk in the academic literature on capacity control. Only the most relevant literature is mentioned. For a recent review, we refer to Gönsch and Hassler (2014) and the references therein.

The need for considering risk-aversion in capacity control was first raised by Lancaster (2003) who proposed a risk-adjusted revenue per available seat mile. Weatherford (2004) then modified the famous EMSR-b heuristic of Belobaba (1992) by substituting revenues with a risk-averse utility function. Barz (2007), Barz and Waldmann (2007), and Feng and Xiao (2008) use an exponential utility function to model risk-aversion, but instead of altering a heuristic, they work with the original DP formulation. Assuming independent demand and a single resource, they show that several well-known properties regarding the structure of an optimal policy carry over from the risk-neutral case. In addition, Barz (2007) extends this analysis to the case of customer choice behavior. Zhuang and Li (2011) examine optimal booking limits with an atemporal utility function to address risk-aversion. Furthermore, there are two publications from Koenig and Meissner (Koenig and Meissner (2015b, 2015c)) who consider target percentile risk and value-at-risk.

Most relevant to our work are Huang and Chang (2011) and Koenig and Meissner (2015a). Similar to our work, they modify existing capacity control approaches to address risk-aversion. In particular, Huang and Chang (2011) heuristically consider risk-aversion via a discount factor on the opportunity cost in the DP formulation. This factor is either constant or a function of remaining demand and capacity. Koenig and Meissner (2015a) extend this analysis. In addition, they consider a discount factor on the opportunity cost from a risk-neutral DP formulation and an alternative function of demand and capacity. However, the approaches are restricted to a few parameters that are calibrated manually. Moreover, as in all the literature on risk-averse capacity control so far, only single-leg settings are considered.

Finally, the literature on risk-averse dynamic pricing is related to us (see Gönsch et al. (2016) for a recent review). The main difference between dynamic pricing and capacity control is that the decision maker influences demand by setting the prices of products instead of choosing the offer set, while the general setting is quite similar (see, e.g., Gallego and van Ryzin (1997) and Talluri and van Ryzin (2004b) for problems with risk-neutral decision makers). Thus, the incorporation of risk-aversion is done in a similar fashion (see, e.g., Li and Zhuang (2009) for utility functions; Feng and Xiao (1999) for revenue variance; Levin et al. (2008) for target percentile risk; Gönsch et al. (2015) for conditional value-at-risk).

2.3 Simulation-based optimization for capacity control

Until now, SBO has only been used in risk-neutral capacity control. For a general overview of SBO please refer to, for example, Gosavi (2015) or Spall (2003). Robinson (1995) was the first to use SBO in the context of revenue management to approximate the optimal booking limit policy in the single-leg case. More recent research derives stochastic gradients of the value function and uses estimates of these gradients in the optimization step. Bertsimas and de Boer (2005) present an algorithm for the improvement of booking limits, which uses a discretization of the state space for value function estimation. Gosavi et al. (2007) show that an algorithm based on simultaneous perturbation for the improvement of booking limits outperforms both EMSR-b and DAVN-EMSR-b for single-leg and network problems, respectively. Topaloglu (2008) and van Ryzin and Vulcano (2008b) improve bid prices and nested protection limits, respectively, by using a continuous approximation of the discrete problem which enables an exact recursive computation of gradients. While all the aforementioned papers follow the independent demand assumption, van Ryzin and Vulcano (2008a) use a procedure similar to van Ryzin and Vulcano (2008b) in order to improve nested protection limits under customer choice behavior. Chaneton and Vulcano (2011) present a stochastic gradient algorithm for improvement of bid prices with customer choice.

However, the use of stochastic gradients is feasible only if a recursive formulation of the value function is available. Unfortunately, this is often not the case or very challenging for the objective functions considered in risk-averse revenue management. Therefore, we concentrate on purely numerical approaches in our paper that can be adapted easily to different objectives. In this sense, Klein (2007) is closest to us. He introduces auto-adaptive bid prices by means of the metaheuristic scatter search assuming independent demand.

3 Risk-averse capacity control using SBO

In this section, we first present an overview of the new framework allowing the incorporation of risk-aversion. Then, we turn to the most important components and describe in detail the risk measures and capacity control approaches considered in this study. Table 2 summarizes the additional notation introduced in the following section.

n^{calib}	number of calibration streams	$U_\gamma(\cdot)$	exponential utility function with level of risk-aversion γ
n^{eval}	number of evaluation streams	$CVaR_\alpha(\cdot)$	conditional value-at-risk (CVaR) at probability level α
R	revenue obtained	θ	(arbitrary) parameters to integrate risk-aversion into capacity control mechanisms
$U(\cdot)$	utility function		
$F(y)$	distribution of total revenue R , i.e., $F(y) = \mathbb{P}(R \leq y)$		

Table 2: Notation introduced in this section

3.1 Overview

Our basic idea is to modify standard approaches appropriately to account for risk-aversion. This modification is governed by parameters θ , which are determined by an out-of-the-box iterative SBO algorithm before the beginning of the booking horizon. The whole process consists of three steps (see Figure 1).

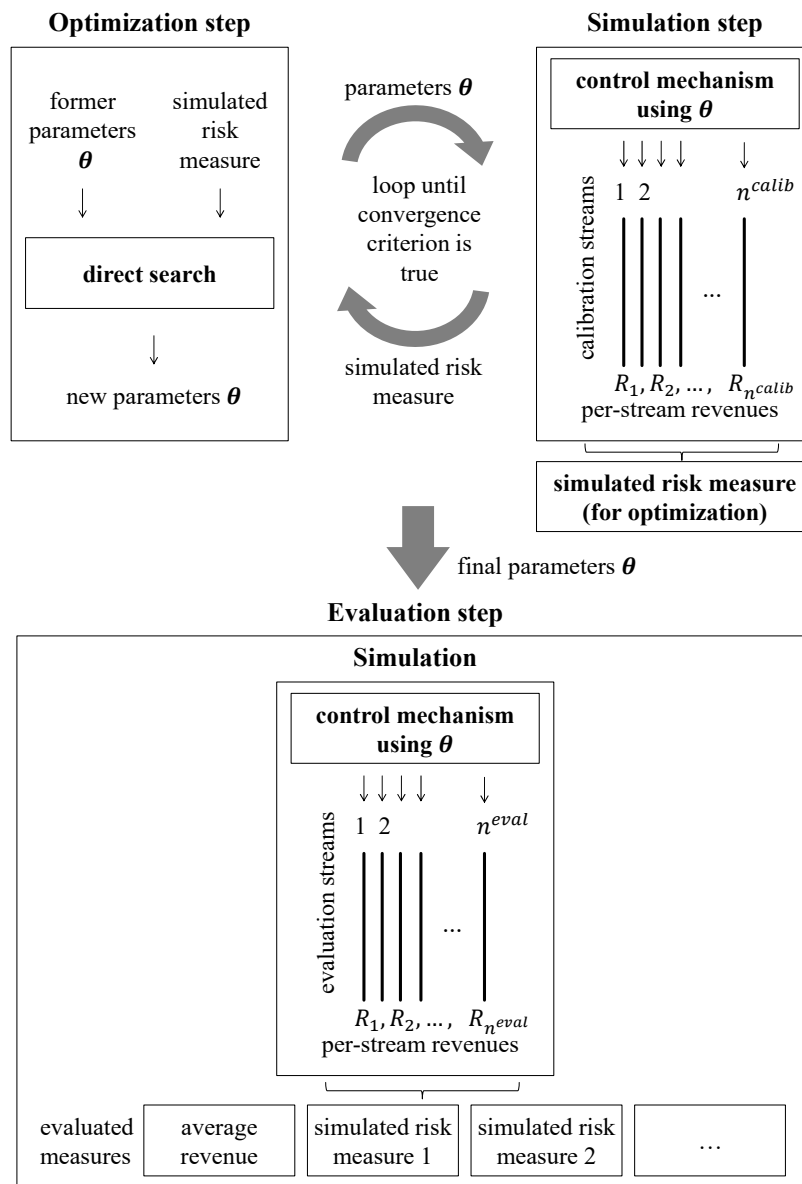


Figure 1: Framework for risk-averse capacity control

The optimization step aims at improving values for the parameters θ . It passes tentative values to the simulation step to estimate their performance. The simulation step in turn mimics sales processes using n^{calib} independent customer demand streams, each encompassing the whole booking horizon. This calibration set is generated in advance according to the firm’s belief about future demand. For each demand stream, the control mechanism with the current values of θ is applied and a total per-stream revenue is obtained. All per-stream revenues are used to calculate a risk measure, which is passed back to the optimization step as an estimate of the θ -values’ performance. Using this

new estimate as well as data from previous iterations, a standard (derivative-free) direct search optimization technique computes new values of θ . These new values are passed to the simulation step again and a new iteration starts. The cycle ends when a predefined convergence criterion is satisfied. The final values of the parameters θ are tested in the evaluation step. Analogously to the simulation step, the resulting control mechanism is applied to n^{eval} demand streams of the evaluation set and various risk measures are calculated. The evaluation step is completely analogous to the simulation step, except that the evaluation set must obviously be independent from the calibration set.

The framework described above can be easily tailored to specific applications by changing two key components that are technically independent from each other: the modified capacity control approach and the optimized risk measure. Accordingly, we will identify the method used with an abbreviation of the form *SBO-[MECHANISM]-[RISKMEASURE]*. In the following, we describe the variants of each of these components we consider in this study. Note that, in addition, the SBO technique can also be varied, but we do not investigate this rather technical issue and stick to a standard approach.

3.2 Risk measures

In the following, we briefly restate the risk measures used in this study. As customers' arrivals and choices are stochastic, total revenue obtained with a given control mechanism is random and denoted by the random variable R with distribution function $F(y) = \mathbb{P}(R \leq y)$. Note that bigger values of R are preferred.

One well established way to address risk-aversion is the use of expected utility which was introduced by von Neumann and Morgenstern (1944). The main idea behind this concept is that decision makers value the same revenue differently due to individual preferences. These preferences are encompassed in an utility function U and two random revenues, say R_1 and R_2 , can be compared by the resulting expected utility, where R_1 is preferred over R_2 if $\mathbb{E}[U(R_1)] \geq \mathbb{E}[U(R_2)]$. Following Barz and Waldmann (2007), we consider an exponential utility function:

$$U_\gamma(R) = 1 - e^{-\gamma \cdot R} \quad (5)$$

The parameter γ indicates the level of risk-aversion. The exponential utility function is the most widely used nonlinear utility function (see, e.g., Corner and Corner (1995)). In our computational study, we abbreviate this risk measure as *Utility*(γ).

The second risk-measure we consider, conditional value-at-risk (CVaR), has attracted a lot of attention over the last decade. For a given probability level $\alpha \in [0,1]$, the CVaR at level α is simply the expectation below the α -quantile of F :

$$CVaR_\alpha(R) = \mathbb{E}[R | R \leq F^{-1}(\alpha)] \quad (6)$$

CVaR is often described as an advancement of the widely popular Value-at-Risk (VaR) to avoid certain theoretical and practical shortcomings of the latter (see, e.g., Artzner et al. (1999)). Note that, to be formally precise, the intuitive definition (6) is valid only for atomless distributions. As revenues are discrete in capacity control, we use CVaR's less intuitive dual representation (not given here; see, e.g., Pflug and Römisch (2007)) to calculate the CVaR. In our computational study, we abbreviate CVaR at level α as *CVaR*(α).

3.3 Capacity control mechanisms

In total, we augment three standard control mechanisms for the optimization of arbitrary risk measures.

The first two mechanisms use bid prices π_{tic_i} that, in case of a single resource, come directly from DP-EV (1) or, in case of multiple resources, from the DP decomposition proposed in Liu and van Ryzin (2008). Thus, the bid prices equal or approximate opportunity cost from the risk-neutral problem. Then, building on Koenig and Meissner (2015a) as well as on Huang and Chang (2011), we integrate a constant factor θ_i to adjust the bid prices to different levels of risk-aversion.

Our first mechanism, *BPF* (“**B**id **P**rice control with **F**actor”), follows a traditional, independent demand bid price control approach. Hence, a product j is available for sale if

$$r_j \geq \sum_{i=1}^m a_{ij} \cdot \theta_i \cdot \pi_{t+1,ic_i}, \quad (7)$$

where the parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_m)^T > \mathbf{0}$ are determined using SBO as described above.

In the second mechanism, *AOF* (“Assortment Optimization with Factor”), we adjust the bid prices within the exact assortment optimization problem. Compared to (7), this approach is able to consider more combinations of products. Accordingly, the offer set is determined by solving

$$\max_{\mathbf{x}} \left\{ \sum_{j=1}^n p_{tj}(\mathbf{x}) \cdot \left(r_j - \sum_{i=1}^m a_{ij} \cdot \theta_i \cdot \pi_{t+1,ic_i} \right) \right\}. \quad (8)$$

In order to solve (8) efficiently during our simulation step, we use the greedy algorithm of Miranda Bront et al. (2009). Although this approach is heuristic in nature, it is known to yield high-quality solutions. Please note that the bid prices (i.e., $\theta_i \cdot \pi_{tic_i}$) used in these approaches are artificially set to infinity if $c_i = 0$. Furthermore, they are state-dependent with regard to the state definition from the risk-neutral problem and represent input parameters to the SBO-algorithm altering the bid price control via the choice of $\boldsymbol{\theta}$.

For our third approach, *BPB* (“Bid Price control with Basis functions”), we broadly follow Klein (2007) and use state-dependent bid prices in (7). The state-dependency is given by a linear model of basis functions:

$$\pi_{tic_i} := \pi_i^0 - \theta_i^{cap} \cdot \frac{c_i}{c_i^0} + \theta_i^{time} \cdot \frac{(T-t+1)}{T} \quad (9)$$

with $\pi_{tic_i} = \infty$ if $c_i = 0$. Again, the parameters $\boldsymbol{\theta}^{cap} = (\theta_1^{cap}, \dots, \theta_m^{cap})^T$ and $\boldsymbol{\theta}^{time} = (\theta_1^{time}, \dots, \theta_m^{time})^T$ are estimated by SBO. The variables c_i and $(T - t + 1)$ sufficiently describe the current booking situation. Note that we normalize these variables to ease the calibration. π_i^0 is our starting bid price coming from a linear approximation of (DP-EV) such as the well-known deterministic linear program (DLP; see, e.g., Talluri and van Ryzin (1998)) or choice-based deterministic linear program (CDLP; see, e.g., Gallego et al. (2004) and Gallego et al. (2004); Liu and van Ryzin (2008)).

4 Simulation study

In this section, we illustrate the impact of our approaches for a risk-averse decision maker, that is, the improvement in risk measure and, if at all, the associated loss in expected revenue. We use four examples that are—as usual in the literature—expressed in airline terminology. However, the results can be transferred to other areas of application. Wherever available, we use standard example networks from literature.

All algorithms were implemented in MATLAB (Version 8, Release R2014b). Linear Programs were solved by the function `linprog` from the Optimization Toolbox, Mixed-Integer Linear Programs by CPLEX from IBM ILOG (Version 12.6). In the optimization step, we used the function `patternsearch` with standard settings from the Global Optimization Toolbox. For each problem instance, the size of the evaluation set is $n^{eval} = 10,000$. Regarding the three SBO-based approaches presented in the previous section, we use a calibration set of $n^{calib} = 5,000$ demand streams. Additional notation introduced in this section is summarized in Table 3.

Single-leg with independent demand	Parallel flights and one-hub network with choice-based demand
p_{tj} probability of selling product j	l customer segment
Single-leg with choice-based demand	\mathcal{C}_l consideration set of segment l
v_{tj} preference weight of product j in period t	λ_l arrival probability of a customer from segment l
v_0 no-purchase preference weight	z_{lj} binary variable indicating whether consideration set \mathcal{C}_l contains product j
	$\mathbf{v}_l = (v_{lj})_{ C_l \times 1}$ preference weights of segment l
	v_{l0} no-purchase preference weight of segment

Table 3: Notation introduced in this section

4.1 Example 1: Small single-leg flight with independent demand

In our first experiment, we consider the classical single-leg example of Lee and Hersh (1993) which was also used by several previous studies on risk-averse capacity control (see, e.g., Barz (2007), Barz and Waldmann (2007), and Koenig and Meissner (2015a)).

It represents a small single-leg flight with a capacity of $c^0 = 10$ seats and $n = 4$ products (booking classes) with revenues $\mathbf{r} = (200, 150, 120, 80)^T$. Demand follows the independent demand assumption, that is, the selling probabilities $p_{tj}(\mathbf{x})$ are independent of $x_i, i \neq j$, and given by

$$p_{tj}(\mathbf{x}) = \begin{cases} p_{tj} & \text{if } x_j = 1 \\ 0 & \text{else} \end{cases}. \quad (10)$$

The booking horizon consists of $T = 30$ periods and is partitioned into five time intervals, so that higher value demand tends to arrive later in the booking horizon (see Table 4).

p_{tj}	$j = 1$	$j = 2$	$j = 3$	$j = 4$
$t = 1, \dots, 5$	0.08	0.08	0.14	0.14
$t = 6, \dots, 12$	0.06	0.06	0.14	0.14
$t = 13, \dots, 19$	0.10	0.10	0.10	0.10
$t = 20, \dots, 26$	0.14	0.14	0.16	0.16
$t = 27, \dots, 30$	0.15	0.15	0	0

Table 4: Purchase probabilities in Example 1

In this subsection, we consider the risk measures CVaR (6) and expected utility with an exponential utility function (5). We combine these risk measures with the control mechanisms *BPF* and *BPB* and, thus, investigate the performance of *SBO-BPF-CVaR*(α), *SBO-BPB-CVaR*(α), *SBO-BPF-Utility*(γ), and *SBO-BPB-Utility*(γ). Because we assume that demand is independent of the offer set, we do not need to consider the capacity control mechanism *AOF*. Furthermore, we implemented the following approaches as benchmarks:

- *BPF*₁ is our benchmark. This is the expected revenue-maximizing policy derived from (1), that is, using decision rule (4) with $\pi_{t1c} := V_t(c) - V_t(c - 1)$.
- *BPF*_{0.8} uses a constant discount factor of 0.8 on the opportunity cost π_{t1c} in line with Koenig and Meissner (2015a) and Huang and Chang (2011).
- *DP-CVaR*(α) is the CVaR-maximizing policy based on the DP formulation of Gönsch and Hassler (2014) and depends on the probability level α .

- $DP\text{-}Utility(\gamma)$ is the expected utility-optimal policy from Barz and Waldmann (2007) and depends on the level of constant absolute risk-aversion γ .

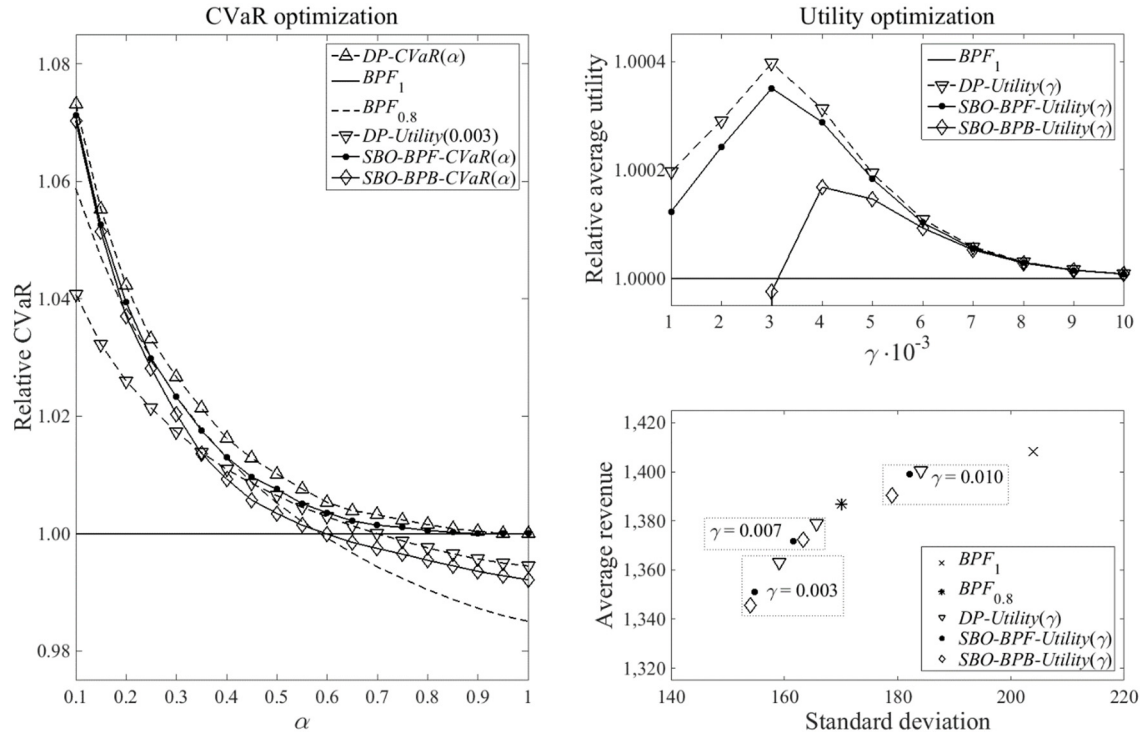


Figure 2: CVaR and average utility in Example 1

In the left (right) part of Figure 2, we consider a CVaR- (utility-) maximizing decision maker and depict the CVaR (utility) relative to that of our benchmark, the expected-value optimal policy from BPF_1 . We calculated and evaluated all policies for $\alpha = 0.1, 0.15, \dots, 1$ ($\gamma = 1 \cdot 10^{-3}, \dots, 10 \cdot 10^{-3}$). Taking a look at the left part of Figure 2, not surprisingly, the DP-based approach $DP\text{-}CVaR(\alpha)$ performs best for all values of α . For $\alpha < 0.6$, all control mechanisms, even $DP\text{-}Utility(0.003)$, considerably improve CVaR in comparison to BPF_1 . A constant discount on opportunity cost ($BPF_{0.8}$), as suggested in previous literature, seems to work very well for $\alpha \in [0.2, 0.4]$ but the results quickly worsen for other values of α . Our simulation-based approach $SBO\text{-}BPF\text{-}CVaR(\alpha)$ is—after $DP\text{-}CVaR(\alpha)$ —the second best control mechanism for all values of α . The factors θ_1 determined by the SBO monotonically increase from 0.4 to 1 in α . This shows the good performance of the intuitively appealing concept of discounts on opportunity cost, which leads to more accepted requests as risk-aversion increases. Moreover, the fact that $SBO\text{-}BPF\text{-}CVaR(\alpha)$ is able to reclaim most of the difference in

CVaR between $BPF_{0.8}$ and $DP-CVaR(\alpha)$ —which is applicable only in single-leg settings due to its inherent DP formulation—is encouraging and underlines the performance of the more general SBO approach. However, $SBO-BPB-CVaR(\alpha)$ and $DP-Utility(0.003)$ yield a poorer performance in this example. This is due to the fact that the linear basis functions of $SBO-BPB-CVaR(\alpha)$ are not able to fully capture the monotonicity of the opportunity cost of the expected revenue-maximizing value function (or, equivalently, the concavity of the value function). Therefore, given such a setting, using a simple discount on opportunity cost is advised. Regarding $DP-Utility(0.003)$, the poorer performance is not surprising as the corresponding policy is optimized in respect to a different risk measure.

Now, please consider the upper right part of Figure 2. The upper bound on average relative utility is given by the exact DP-based approach $DP-Utility(\gamma)$ of Barz and Waldmann (2007). Because the differences in relative average utility between the different control mechanisms are quite small, we chose to limit the range of values to a small interval, thus excluding $BPF_{0.8}$ from the figure due to a poorer performance. $SBO-BPB-Utility(\gamma)$ works fine for maximizing utility, as the results are nearly identical to $DP-Utility(\gamma)$. The factors θ_1 are again discounts that range from 0.67 to 0.9 and decrease with increasing risk-aversion γ . Similar to the optimization of CVaR, the results of $SBO-BPB-Utility(\gamma)$ are slightly worse. However, all control mechanisms, including BPF_1 , show practically identical results for $\gamma \geq 0.007$.

On the lower right part of Figure 2, we compare average revenue and standard deviation for $\gamma \in \{0.003, 0.007, 0.01\}$. Obviously, higher values of γ lead to a smaller average revenue but also a smaller standard deviation of revenues, yielding some kind of efficient frontier. This shows that although the differences in relative utility are often negligible, the approaches lead to different policies.

4.2 Example 2: Single-leg flight with choice-based demand

In the remainder of the paper, we assume customer choice behavior. Hence, our main benchmark mechanism is AOF_1 with its near-optimal policy regarding expected revenue

and—unless stated otherwise—all results are given relative to this benchmark. Moreover, we now focus on the optimization of CVaR. We consider all three SBO-based mechanisms and additionally state the results of BPF_1 . Please note that $DP-CVaR(\alpha)$ is not tractable for the following examples.

Unfortunately, we are not aware of an established choice-based single-leg setting from the literature for capacity control. There are only a few settings complementing analytical results. For example, Talluri and van Ryzin (2004a) use a simple example to illustrate demand estimation by an expectation-maximization algorithm and as a proof of concept for their DP formulation. However, in their example, capacity is not scarce (i.e., opportunity cost equals zero) and the authors only have to solve the same assortment optimization problem over time. Nonetheless, the following example is structurally similar.

In this subsection, we consider a single-leg flight with a capacity of $c^0 = 50$, four products with revenues $\mathbf{r} = (1000, 800, 600, 400)^T$ and $T = 110$ periods. Demand follows a multinomial logit model. Thus, the purchase probabilities $p_{tj}(\mathbf{x})$ depend on product-specific preference weights v_{tj} as well as the no-purchase preference weight $v_0 = 1$ and are given by

$$p_{tj}(\mathbf{x}) = \frac{v_{tj} \cdot x_j}{1 + \sum_{k=1}^n v_{tk} \cdot x_k}. \quad (11)$$

We consider two variants regarding the distribution of demand over time. In the first variant, the purchase probabilities are time-homogenous. In the second variant, higher value demand tends to arrive later in the booking horizon. We call these settings *time-homogenous* and *low-before-high*, respectively. The corresponding values of v_{tj} are given in Table 5. The 10^{-5} values in the second demand variant lead to virtually no demand for the corresponding products, but due to some technicalities, the weights must be strictly positive (also this is often not explicitly stated in the literature).

v_{tj}	$j = 1$	$j = 2$	$j = 3$	$j = 4$
<i>time-homogenous</i>				
$t = 1, \dots, 110$	0.05	0.1	0.5	0.6
<i>low-before-high</i>				
$t = 1, \dots, 80$	10^{-5}	0.2	0.6	0.8
$t = 81, \dots, 110$	0.2	0.3	10^{-5}	10^{-5}

Table 5: Preference weights in Example 2

Figure 3 shows the CVaR of all control mechanisms relative to AOF_1 . Note that the spread of relative CVaR is higher in the *low-before-high*-setting. This reflects a well-known effect of capacity control, namely that the decision problem becomes more challenging when demand tends to arrive in low-before-high order.

Naturally, $SBO-AOF-CVaR(\alpha)$ performs best and, in both examples, with small benefits over $SBO-BPB-CVaR(\alpha)$ as well as larger benefits over $SBO-BPF-CVaR(\alpha)$. Interestingly, even the standard bid price control BPF_1 is often competitive and yields a higher CVaR than AOF_1 for low levels of α because more low value products are offered for sale. However, for medium and high values of α , there are severe losses in CVaR. With SBO, these losses can be successfully reduced.

Regarding θ_1 , we observe values of 0.45 to 1 for $SBO-AOF-CVaR(\alpha)$ that are increasing in α and which represent discounts on the opportunity cost analogously to Example 1. Regarding $SBO-BPF-CVaR(\alpha)$, θ_1 ranges from 0.85 to 1.87 and also increases with α . This reflects that with a bid price control, there is a trade-off between a discount on the opportunity cost to allow for risk-aversion and a markup to prohibit buy-down behavior. This effect is better captured by the state-dependent bid prices of $SBO-BPB-CVaR(\alpha)$ in comparison with $SBO-BPF-CVaR(\alpha)$, which uses a constant markup over the whole booking horizon. In a risk-neutral setting, increased bid prices to induce higher value demand were, for example, also observed in Meissner and Strauss (2012a).




Figure 3: CVaR in Example 2

We now delve deeper into the opportunities and threats that accompany risk-averse capacity control. For all SBO-based mechanisms and three selected values of α , the upper part of Figure 4 shows the absolute gains in CVaR compared to the absolute gains in average revenue. Please note that the gains in CVaR are subject to the specific level of risk-aversion α and, thus, need to be treated with caution. A higher gain in CVaR usually leads to a bigger loss in expected revenue. For example, in *low-before-high*, improving the $CVaR_{0.4}$ by around 250 costs 500 in average revenue (see the upper right part of Figure 4). The lower part of Figure 4 compares the sampled distribution of total revenues of $SBO-AOF-CVaR(0.4)$ and AOF_1 over the evaluation streams. In line with the results from the optimization of utility in Example 1, risk-averse capacity controls leads to a smaller support of the distribution and shorter tails. In other words, extreme outcomes are less likely.




Figure 4: CVaR vs. expected value and revenue distribution in Example 2

4.3 Example 3: Parallel flights with choice-based demand

Our third example is based on the parallel flight network of Miranda Bront et al. (2009). It consists of three flights with two products defined on each flight, that is, one low-class and one high-class product. In this example, we additionally investigate the impact of different capacity provision on the risk-profile: First, we consider an initial capacity of $\mathbf{c}^0 = (27, 45, 36)^T$ and, second, an initial capacity of $\mathbf{c}^0 = (21, 35, 28)^T$. Demand follows a mixture of multinomial logit models. More precisely, customers belong to different market segments $l = 1, \dots, 4$, each of which has a subset of products to consider for purchase, namely the consideration set \mathcal{C}_l . The variable z_{lj} indicates whether product $j \in \mathcal{C}_l$ ($z_{lj} = 1$) or not ($z_{lj} = 0$). A customer from segment l arrives with probability λ_l and has preference weights $\mathbf{v}_l = (v_{lj})_{|\mathcal{C}_l| \times 1}$ as well as v_{l0} for the no-purchase alternative. Note that v_{lj} is only defined if $z_{lj} = 1$. Demand is time-homogenous over the booking horizon of $T = 300$ periods. Then, the probability of selling product j in period t is given by

$$p_{tj}(\mathbf{x}) = \sum_{l=1}^4 \lambda_l \cdot \frac{v_{lj} \cdot z_{lj} \cdot x_j}{v_{l0} + \sum_{k=1}^n v_{lk} \cdot z_{lk} \cdot x_k}. \quad (12)$$

The remaining data is summarized in Table 6 in Appendix A. The corresponding relative CVaRs are shown in Figure 5. Similar to the previous example, the absolute CVaR gains vs. absolute expected revenue gains and the sampled distribution of total revenues for $\alpha = 0.4$ are illustrated in Figure 6.

In the following, we summarize the key observations complementing the former results. First, comparing the results for the two initial capacities, we see that only minor variations of the setting can lead to large differences in the risk profile. Second, there are settings, such as the first initial capacity, where the consideration of risk-aversion is more or less negligible. Given such a setting, the better the expected revenue of a policy is, the better is its CVaR for almost all levels α and vice versa. More precisely, optimizing one can also increase the other and it suffices to optimize expected revenue. Third, standard bid price controls such as BPF_1 can perform quite poorly when considering customer choice behavior (second initial capacity) and SBO can successfully address this. In this instance, $SBO-BPF-CVaR(\alpha)$ uses factors θ_i of up to 3.97 and outperforms the other approaches. It is usually even better than the near-exact assortment optimization in AOF_1 . For example, the gain of $SBO-AOF-CVaR(1)$ in expected revenue over AOF_1 is almost 1%. This remarkable result can only be explained with the fact that all approaches, including AOF_1 , use approximate bid prices from a DP decomposition instead of the exact opportunity cost from the intractable DP formulation.

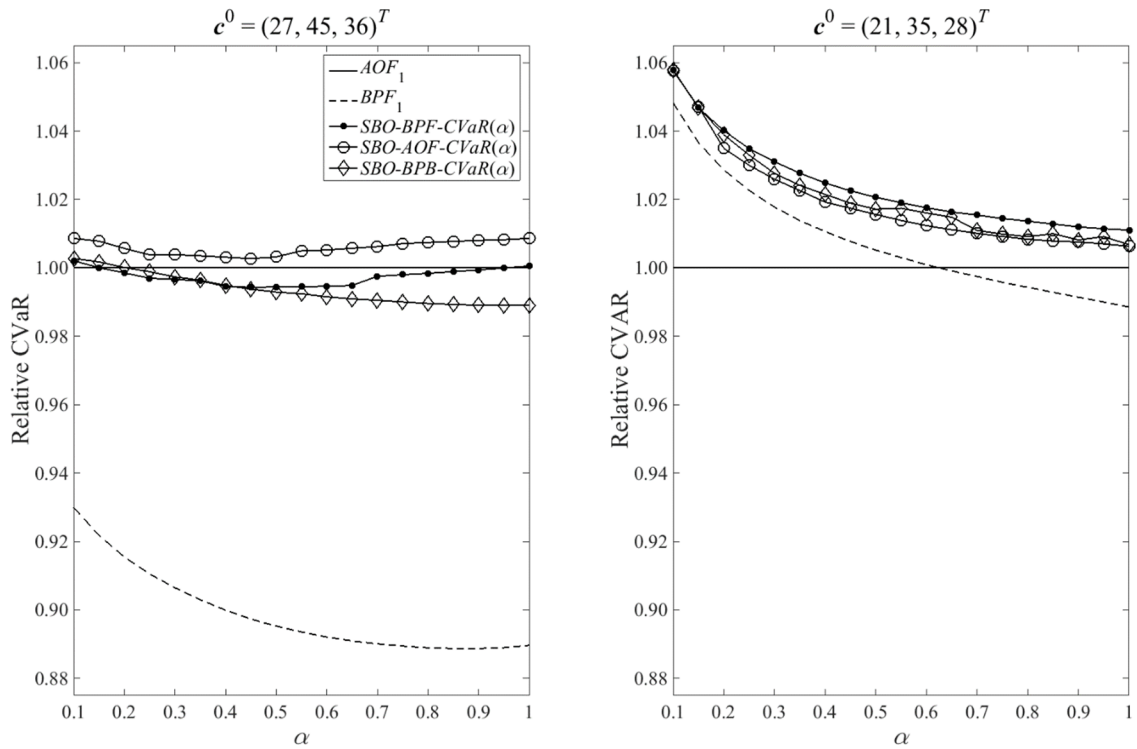


Figure 5: CVaR in Example 3

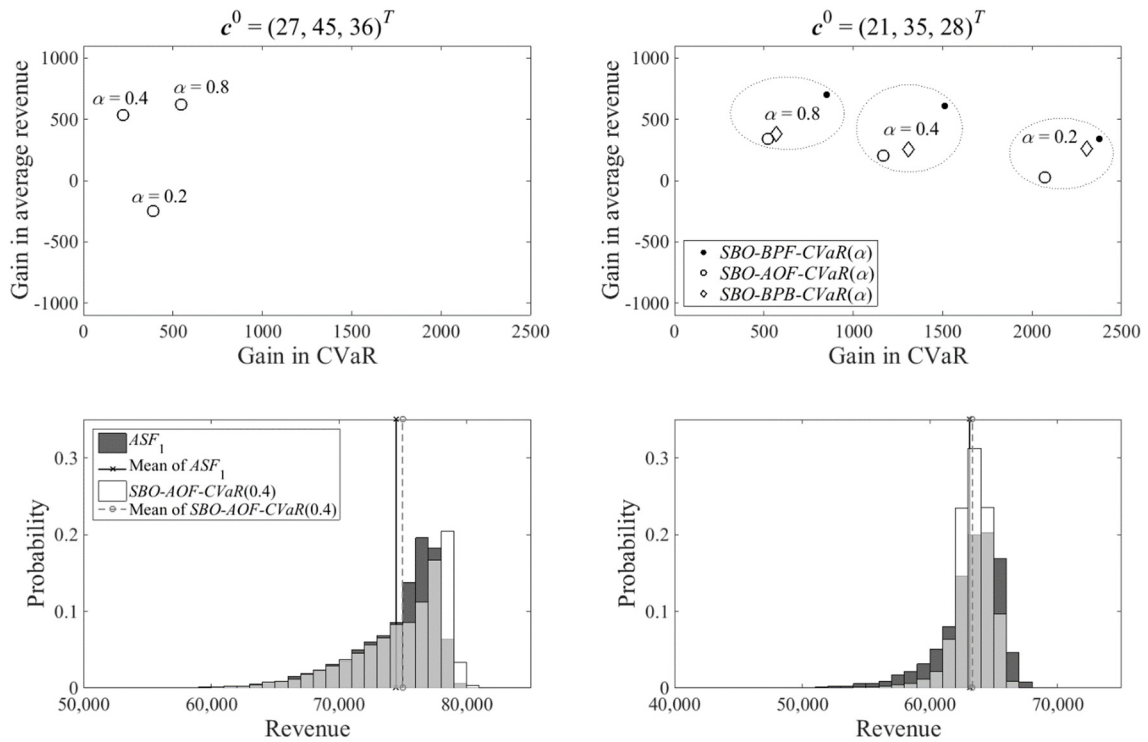


Figure 6: CVaR vs. expected value and revenue distribution in Example 3

4.4 Example 4: One hub-network with choice-based demand

The last example is based on an airline network structure from Meissner and Strauss (2012b) with one hub H connecting two non-hub cities A and B with four flight legs (see Figure 7). There are six itineraries (A to H, A to B via H, H to B, B to H, B to A via H, and H to A). For each itinerary, one high-class and one low-class product are available. The demand behavior is the same as in the parallel flight example. For each itinerary, there is one customer segment with a higher preference for the low-class product. A detailed description of products (revenues r_j and capacity consumptions \mathbf{a}_j) and segments (consideration sets \mathcal{C}_l , preference weights \mathbf{v}_l , no-purchase preference weights v_{l0} , and segment probabilities λ_l) can be found in Table 7 in Appendix A.




Figure 7: One-hub network of Example 4

We assume an initial capacity $c_h^0 = 15$ ($c_h^0 = 60$) for all h and a booking horizon length $T = 300$ ($T = 1200$). The corresponding relative CVaRs are shown in Figure 8.

In general, the observations confirm the results of the previous examples. However, it is remarkable that large gains in CVaR and expected revenue in settings with connecting flights are possible. This is because AOF_1 decomposes the network by flights and, apparently, this does not sufficiently capture the network effects. SBO is able to remedy this shortcoming and, thus, all three SBO-based approaches perform quite well. Comparing the two initial capacities shows that the differences between all mechanisms decline as we scale up the size of the network.




Figure 8: CVaR in Example 4

5 Discussion

After analyzing the numerical results in detail in the previous section, we now take a broader perspective and discuss the relevance of our work.

First, the relevance of risk-aversion in revenue management—and particularly network revenue management—is the foundation of this paper, although risk-neutrality has been taken for granted for a long time in the literature. One reason is probably that risk-neutrality leads to mathematically simpler models. However, many people who are new to revenue management consider this assumption counter-intuitive and many industry partners question it at first.

Lancaster (2003) was the first to raise the issue of risk in revenue management. He pointed out that airlines, like all businesses, face risks which should be managed appropriately. By contrast, he observes that revenue management considers only the reward side, that is, increasing expected revenue and completely ignores the risks assumed in doing so. Despite this early work, most authors cite experiences from practice to show

the relevance of risk-aversion. For example, smaller airlines asked a consultant about risk-averse capacity control (see Weatherford (2004)). Two other examples are due to Levin et al. (2008). Event promoters may organize only a few large events per year in locations that are very expensive to rent. Accordingly, their first priority is to recover this fixed cost (see Levin et al. (2008)). In other industries, a manager's primary concern is often to provide stable results because negative news can lead to negative stock market assessments that can far outweigh the marginal revenue advantages of a risk-neutral policy.

In contrast to the above-mentioned rather small businesses, many companies have a large number of similar events. Thus, the law of large numbers ensures that the average revenue of each event is maximized and also quite stable when using a risk-neutral model that focuses on the expected revenue. For example, network airlines have hundreds, major ones even several thousands of take-offs every day. Although risk-neutrality may be appropriate for these companies as a whole, it may not be appropriate for every department and individual decision maker, leading to missing acceptance of risk-neutral revenue management systems. For example, a consultant's clients were not comfortable with their risk-neutral revenue management system (see Barz (2007)). They manually altered the forecast to obtain less aggressive (and risky) results. Singh (2011) observed that analysts' individual risk-aversion has a huge impact on their decisions when overwriting a revenue management system's output at a cruise line company. He attributes this behavior largely to their personality because they made decisions about exactly the same issues and possessed identical information.

Second, we would like to point out that, regarding the company as a whole, the need to incorporate risk-aversion declines with the network size due to compensatory effects. However, our examples showed that there are also bigger settings where risk-aversion is relevant in the sense that a risk-averse solution differs from a risk-neutral one. For this purpose, we used established examples from literature both for single-leg and network capacity control with mid-sized capacity. In our opinion, these networks serve as a good representation of the sub-networks that individual decision makers may control. None-

theless, it would be worthy of future research to investigate the impact of risk-averse control mechanisms in large-scale problem instances from practice.

Third, our results show that we were able to sufficiently address customer choice behavior in most cases by using bid price rules instead of solving the exact assortment optimization problem. This is in line with the literature on the optimization of expected revenue: Chaneton and Vulcano (2011) and Meissner and Strauss (2012a) make similar observations. But, despite of their widespread use, bid price controls can sometimes yield a comparatively poor performance. This may be due to the fact that bid price controls are not always able to represent the optimal policy in networks, especially if customer choice is considered (see, e.g., Talluri and van Ryzin (1998)). By contrast, the solution of the assortment optimization problem is able to represent all decision options. Thus, $SBO-AOF-CVaR(\alpha)$ performs slightly better than the bid price controls. Nonetheless, our results indicate that $SBP-BPF-CVaR(\alpha)$ and $SBO-BPB-CVaR(\alpha)$ perform very well and, thus, explain the favoritism of bid price controls in practice due to the trade-off between solution quality and simplicity.

To summarize the discussion so far, the overall framework works quite well because arbitrary controls may be designed and optimized regarding a certain risk measure. We focused on enhancing existing control mechanisms with risk-averse components, since commercial revenue management systems are fixed in the long run. Given a control mechanism, our framework always improves the results of the original control.

Finally, the framework presented can also be used to capture risk-aversion in dynamic pricing, where a firm decides on the products' prices instead of their availability. Thus, instead of the assortment optimization problem (3), the firm has to solve a pricing problem in each period to determine the products' prices. As in capacity control, the value of future sales is reflected by opportunity cost which is often approximated by bid prices. Therefore, our framework can be easily applied to risk-averse dynamic pricing because the bid prices can be modified via tunable parameters that capture risk-aversion. Moreover, if continuous prices are allowed, the pricing problem can be solved analytically for many demand models and, thus, much faster than the assortment optimization prob-

lem we consider. This could allow the SBO to perform more simulation runs, possibly leading to even better results.

6 Conclusion

We presented a flexible and modular framework for risk-averse capacity control that offers several advantages compared to existing approaches. First, the practical decision rules we consider can be implemented easily in existing operational systems because they build on well-established standard risk-neutral control mechanisms. Second, using SBO, the control mechanisms can be tailored to every risk measure. Third, because SBO-algorithms are meanwhile widely available in standard software, the only prerequisite for using this model-free framework is being able to undertake Monte-Carlo simulations of the arrival process and choice behavior of customers. There is no need for a DP formulation of the decision problem, which is prohibitive for most risk measures, or an explicit model of customer behavior.

For demonstration purposes, we conducted a simulation study with the widely used multinomial logit model, but our approach admits the use of any other choice model. Based on CVaR and expected utility, we showed that small and intuitive modifications in standard control mechanisms, if designed properly, can be sufficient to successfully incorporate risk-aversion into capacity control, including network settings and customer choice. This usually leads to a narrower distribution of revenues in comparison to standard controls as well as more predictable and stable revenues.

Appendix

A Product and segment data in Example 3 and 4

Regarding Example 3 (parallel flights) and Example 4 (one hub), Table 6 and Table 7 summarize the remaining product and segment data.

Products			Segments				
j	\mathbf{a}_j	r_j	l	\mathcal{C}_l	\mathbf{v}_l	v_{l0}	λ_l
1	$(1, 0, 0)^T$	400	1	{2, 4, 6}	$(5, 10, 1)^T$	1	0.1
2	$(1, 0, 0)^T$	800	2	{1, 3, 5}	$(5, 1, 10)^T$	5	0.15
3	$(0, 1, 0)^T$	500	3	{1, 2, 3, 4, 5, 6}	$(10, 8, 6, 4, 3, 1)^T$	5	0.2
4	$(0, 1, 0)^T$	1000	4	{1, 2, 3, 4, 5, 6}	$(8, 10, 4, 6, 1, 3)^T$	1	0.05
5	$(0, 0, 1)^T$	300					
6	$(0, 0, 1)^T$	600					

Table 6: Product and segment description in Example 3

Products			Segments				
j	\mathbf{a}_j	r_j	l	\mathcal{C}_l	\mathbf{v}_l	v_{l0}	λ_l
1	$(1, 0, 0, 0)^T$	300	1	{1,2}	$(0.5, 2)^T$	1	0.1
2	$(1, 0, 0, 0)^T$	150	2	{3,4}	$(0.5, 2)^T$	1	0.06
3	$(1, 1, 0, 0)^T$	600	3	{5,6}	$(0.5, 2)^T$	1	0.1
4	$(1, 1, 0, 0)^T$	300	4	{7,8}	$(0.5, 2)^T$	1	0.1
5	$(0, 1, 0, 0)^T$	350	5	{9,10}	$(0.5, 2)^T$	1	0.09
6	$(0, 1, 0, 0)^T$	175	6	{11,12}	$(0.5, 2)^T$	1	0.07
7	$(0, 0, 1, 0)^T$	300					
8	$(0, 0, 1, 0)^T$	150					
9	$(0, 0, 1, 1)^T$	500					
10	$(0, 0, 1, 1)^T$	250					
11	$(0, 0, 0, 1)^T$	250					
12	$(0, 0, 0, 1)^T$	125					

Table 7: Product and segment description in Example 4

References

- Artzner P, Delbaen F, Eber J, Heath D (1999) Coherent measures of risk. *Mathematical Finance* 9(3): 203–228.
- Barz C (2007) Risk-averse capacity control in revenue management. Springer, Berlin.
- Barz C, Waldmann KH (2007) Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research* 65(3): 565–579.
- Belobaba PP (1992) Optimal vs. heuristic methods for nested seat allocation. Presentation at AGIFORS Reservations and Yield Management Study Group, Brussels.
- Bertsimas D, de Boer S (2005) Simulation-based booking limits for airline revenue management. *Operations Research* 53(1): 90–106.
- Chaneton JM, Vulcano G (2011) Computing bid prices for revenue management under customer choice behavior. *Manufacturing & Service Operations Management* 13(4): 452–470.
- Corner JL, Corner PD (1995) Characteristics of decisions in decision analysis practice. *Journal of the Operational Research Society* 46(3): 304–314.
- Feng Y, Xiao B (1999) Maximizing revenues of perishable assets with a risk factor. *Operations Research* 47(2): 337–341.
- Feng Y, Xiao B (2008) A risk-sensitive model for managing perishable products. *Operations Research* 56(5): 1305–1311.
- Gallego G, Iyengar G, Phillips RL, Dubey A (2004) Managing flexible products on a network. Working paper, Columbia University, New York.
- Gallego G, van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* 45(1): 24–41.
- Gönsch J, Hassler M (2014) Optimizing the conditional value-at-risk in revenue management. *Review of Managerial Science* 8(4): 495–521.
- Gönsch J, Hassler M, Schur R (2016) Optimizing conditional value-at-risk in dynamic pricing. Working paper, University of Augsburg, Augsburg.
- Gosavi A (2015) Simulation-based optimization: Parametric optimization techniques and reinforcement. 2nd ed. Springer, New York.
- Gosavi A, Ozkaya E, Kahraman AF (2007) Simulation optimization for revenue management of airlines with cancellations and overbooking. *OR Spectrum* 29(1): 21–38.
- Huang K, Chang K (2011) A model for airline seat control considering revenue uncertainty and risk. *Journal of Revenue and Pricing Management* 10(2): 161–171.
- Klein R (2007) Network capacity control using self-adjusting bid-prices. *OR Spectrum* 29(1): 39–60.
- Koenig M, Meissner J (2015a) Risk management policies for dynamic capacity control. *Computers & Operations Research* 59(2015): 104–118.

- Koenig M, Meissner J (2015b) Risk minimising strategies for revenue management problems with target values. *Journal of the Operational Research Society* 67(3): 402–411.
- Koenig M, Meissner J (2015c) Value-at-risk optimal policies for revenue management problems. *International Journal of Production Economics* 166(2015): 11–19.
- Lancaster J (2003) The financial risk of airline revenue management. *Journal of Revenue and Pricing Management* 2(2): 158–165.
- Lee TC, Hersh M (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27(3): 252–265.
- Levin Y, McGill J, Nediak M (2008) Risk in revenue management and dynamic pricing. *Operations Research* 56(2): 326–343.
- Li MZ, Zhuang W (2009) Risk-sensitive dynamic pricing for a single perishable product. *Operations Research Letters* 37(5): 327–332.
- Liu Q, van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2): 288–310.
- Meissner J, Strauss A (2012a) Improved bid prices for choice-based network revenue management. *European Journal of Operational Research* 217(2): 417–427.
- Meissner J, Strauss A (2012b) Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research* 216(2): 459–468.
- Miranda Bront JJ, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. *Operations Research* 57(3): 769–784.
- Pflug G, Römisches W (2007) *Modeling, measuring and managing risk*. World Scientific, Singapore.
- Phillips RL (2005) *Pricing and revenue optimization*. Stanford University Press, Stanford.
- Robinson LW (1995) Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research* 43(2): 252–263.
- Singh P (2011) Risk in revenue management: Acknowledging risk’s existence and knowing how to minimize it. *Inform Analytics Magazine* May/June: 16–19.
- Spall JC (2003) *Introduction to stochastic search and optimization: Estimation, simulation, and control*. Wiley, Hoboken, New Jersey.
- Talluri K, van Ryzin G (1998) An analysis of bid-price controls for network revenue management. *Management Science* 44(11): 1577–1593.
- Talluri KT, van Ryzin G (2004a) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.

- Talluri KT, van Ryzin G (2004b) The theory and practice of revenue management. Springer, New York.
- Topaloglu H (2008) A stochastic approximation method to compute bid prices in network revenue management problems. *INFORMS Journal on Computing* 20(4): 596–610.
- van Ryzin G, Vulcano G (2008a) Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing & Service Operations Management* 10(3): 448–467.
- van Ryzin G, Vulcano G (2008b) Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research* 56(4): 865–880.
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton University Press, Princeton.
- Weatherford LR (2004) EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management* 3(3): 277–284.
- Zhang D, Adelman D (2009) An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Science* 43(3): 381–394.
- Zhuang W, Li MZF (2011) Revenue optimization of risk-averse managers with atemporal utility. *Journal of Revenue and Pricing Management* 10(5): 424–437.

III Fazit und Ausblick

Gegenstand der vorliegenden Dissertationsschrift war die (heuristische) Lösung praxisrelevanter Entscheidungsprobleme im Revenue Management. Ein konkreter Anwendungsfall bestand zum einen in der Berücksichtigung nachfrageseitiger Substitution in Form von Kundenwahlverhalten (vgl. Beiträge B1, B2 und B5). Dies wirkt sich im Vergleich zum klassischen Revenue Management vor allem auf die Komplexität des Entscheidungsproblems, das in jedem Zustand des zugrundeliegenden dynamischen Optimierungsproblems zu lösen ist, aus. Zum anderen wurden mit der Berücksichtigung anbieterseitiger Substitution (vgl. Beiträge B2, B3 und B4) sowie der Annahme eines risikoaversen Entscheiders (vgl. Beitrag B5) zwei Anwendungsfälle betrachtet, bei denen sich die grundlegende Definition des Zustandsraums ändert, so dass Standardmethoden des Revenue Managements nicht ohne Weiteres verwendet werden konnten.

Methodisch lag der Fokus deshalb vor allem auf der Untersuchung simulationsbasierter approximativer, dynamischer Optimierung im engeren Sinne (vgl. Beitrag B1) sowie simulationsbasierter Optimierung zur Verbesserung existierender Standardmethoden (vgl. Beiträge B3 und B5). In beiden Fällen werden – im Gegensatz zu den in der Literatur verbreiteten, überwiegend analytisch motivierten Ansätzen – Simulationspfade der echten Nachfrage verwendet, um im Laufe von Iterationen eine gute Politik zu lernen.

Ein wesentlicher Vorteil eines solchen Vorgehens ist, dass die Bestimmung einer guten Politik unter Zuhilfenahme von Standardsoftware erfolgen kann. Sowohl Solver für Kleinste-Quadrate-Probleme, wie sie bei der Schätzung von Approximationen einer Wertfunktion zum Einsatz kamen, als auch Metaheuristiken zur Kalibrierung weniger Parameter im Rahmen einer simulationsbasierten Optimierung sind bereits vielfach in Software zur Entscheidungsunterstützung integriert. In Rahmen dieser Dissertationsschrift wurde dabei auf in MATLAB verfügbare Routinen sowie auf CPLEX von IBM ILOG zurückgegriffen.

Ein zweiter, entscheidender Vorteil ist, dass vergleichsweise wenige Annahmen bzgl. des Kundenwahlverhaltens getroffen werden müssen. Während in der Literatur zum

Revenue Management aufgrund seiner einfachen Handhabbarkeit nahezu ausschließlich das MNL-Modell untersucht wird, erfordern die vorgeschlagenen Methoden lediglich die Verfügbarkeit von Simulationspfaden (vgl. Powell (2011), Kap. 4.2.3), die bspw. mittels Monte-Carlo-Simulation erzeugt (vgl. Domschke et al. (2015), Kap. 10.1.1) für eine Einführung in Monte-Carlo-Simulation) oder Realweltdaten in Form historischer Buchungen sein können. Damit können auch fortgeschrittene Formen von Kundenwahlverhalten untersucht werden, bspw. falls Kunden komplexere Optimierungsprobleme zur Produktauswahl lösen (vgl. Mayer und Steinhardt (2016)) oder heuristische Entscheidungsregeln befolgen (vgl. Gigerenzer und Gaissmaier (2011)).

Insgesamt konnte gezeigt werden, dass sämtliche vorgeschlagenen Lösungsmethoden für viele der in der Literatur existierenden Standard-Probleminstanzen statistisch signifikant höhere durchschnittliche Erlöse liefern. In diesem Zusammenhang sei erwähnt, dass im Revenue Management bereits kleine Erlössteigerungen aufgrund ihrer Hebelwirkung einen erheblichen Einfluss auf das Gesamtergebnis des Unternehmens haben.

Zudem wurden in dieser Dissertationsschrift einige Ansätze zu einer besseren Ausnutzung der zeitlichen Flexibilität, wie sie häufig mit flexiblen Produkten und Upgrades einhergeht, aufgezeigt. Dennoch steht die Forschung zu (weiteren Formen) anbieterseitiger Substitution immer noch am Anfang. Es ist wohl weiterhin mit dem Verkauf innovativer Produkte in der Praxis und in Folge mit der Untersuchung der Vorteilhaftigkeit dieser Produkte in der Forschung zu rechnen. Beispielsweise ist ein vermehrter Einsatz kostenpflichtiger Upsells anstelle kostenloser Upgrades zu beobachten. Anbieter von Software-Lösungen finden sich unter anderem in der Passagierluftfahrt und der Hotellerie, bspw. Optiontown (www.optiontown.com) und nor1 (www.nor1.com). Ähnlich wie bei flexiblen Produkten und Upgrades lassen sich vermutlich auch Entscheidungsprobleme des Revenue Managements bei Berücksichtigung solcher Upsells nicht mehr über einen rein ressourcenbasierten Zustandsraum beschreiben. Da für derartige, vergleichsweise komplexe Problemstellungen ein Einsatz herkömmlicher, analytisch motivierter Methoden fraglich erscheint, liegen Anpassungen und Weiterentwicklungen der in dieser Dissertationsschrift vorgeschlagenen Lösungsmethoden nahe.

Literaturverzeichnis

- Domschke W, Drexl A, Klein R, Scholl A (2015) Einführung in Operations Research. 9. Aufl. Springer, Berlin.
- Gigerenzer G, Gaissmaier W (2011) Heuristic decision making. *Annual review of psychology* 62: 451–482.
- Mayer S, Steinhardt C (2016) Optimal product line pricing in the presence of budget-constrained consumers. *European Journal of Operational Research* 248(1): 219–233.
- Powell WB (2011) *Approximate dynamic programming: Solving the curses of dimensionality*. 2. Aufl. Wiley, Hoboken, New Jersey.