Εθνικο Μετσοβιο Πολυτεχνειο
Διατμηματικο Μεταπτυχιακο Προγραμμα Σπουδων
Επιστημη Δεδομενων και Μηχανικη Μαθηση

On finding EFX allocations: conditional approximations and tiered rankings

Μεταπτυχιακη Διπλωματικη Εργασια

του

ΣΑΝΤΟΡΙΝΑΙΟΥ Γ. ΧΡΙΣΤΟΔΟΥΛΟΥ

Επιβλέπων: Παγουρτζής Αριστείδης
Καθηγητής Ε.Μ.Π.

Εργαστηριο Λογικης και Επιστημης Υπολογιστων
Αθήνα, Ιούλιος 2022

Εθνικό Μετσόβιο Πολυτεχνείο
Διατμηματικό Μεταπτυχιακό Πρόγραμμα Σπουδών
Επιστήμη Δεδομένων και Μηχανική Μάθηση
Εργαστήριο Λογικής και Επιστήμης Υπολογιστών

On finding EFX allocations: conditional approximations and tiered rankings

Μεταπτυχιακή Διπλωματική Εργασία

του

ΣΑΝΤΟΡΙΝΑΙΟΥ Γ. ΧΡΙΣΤΟΔΟΥΛΟΥ

Επιβλέπων: Παγουρτζής Αριστείδης
Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 4η Ιουλίου 2022.

(Υπογραφή)	*(Υπογραφή)*	*(Υπογραφή)*
....	....	....
Παγουρτζής Αριστείδης	Μαρκάκης Ευάγγελος	Φωτάκης Δημήτριος
Καθηγητής Ε.Μ.Π.	Αν. Καθηγητής Ο.Π.Α.	Αν. Καθηγητής Ε.Μ.Π.

Αθήνα, Ιούλιος 2022

(Υπογραφή)

....

Σαντορινaιοσ Γ. Χριστοδογλοσ
MSc Επιστήμης Δεδομένων και Μηχανικής Μάθησης Ε.Μ.Π.

Ευχαριστίες

Αρχικά, θα ήθελα να ευχαριστήσω τον κ.Μαρκάκη που μου έδωσε την δυνατότητα να εργαστώ πάνω σε μία δραστήρια ερευνητική κατεύθυνση και να παράξω και ορισμένα αποτελέσματα. Φυσικά, χωρίς την συμβολή του και τις συχνές μας συζητήσεις ο όγκος αυτών των αποτελεσμάτων θα ήταν αισθητά μικρότερος.

Επίσης, ευχαριστώ ιδιαίτερα και τα μέλη της κριτικής επιτροπής, τον κ. Παγουρτζή και τον κ. Φωτάκη, τόσο για τη συμμετοχή τους σε αυτή όσο και για τα μαθήματά τους αυτά τα δύο χρόνια του μεταπτυχιακού.

Τέλος, όπως πάντα, θα ήθελα να ευχαριστήσω την οικογένειά μου.

Περίληψη

Σε αυτή τη διπλωματική εργασία, μελετάμε το πρόβλημα του διακριτού δίκαιου διαμοιρασμού, δηλαδή της ανάθεσης αδιαίρετων αγαθών σε παίκτες με δίκαιο τρόπο. Το πρόβλημα πηγάζει από πλήθος εφαρμογών, από το μοίρασμα παιχνιδιών σε παιδιά μέχρι το χώρισμα κληρονομιών, όπου συνιδιοκτησία αγαθών ή χρηματικά ανταλλάγματα δεν επιτρέπονται. Δεδομένου ότι η δικαιοσύνη είναι μια δύσκολη ποσοτικοποιήσιμη ιδέα, πολλές έννοιες έχουν αναπτυχθεί στο πέρασμα των χρόνων. Δυστυχώς, η παρουσία αδιαίρετων αγαθών τις καθιστά μη εφικτές. Για την αντιμετώπιση αυτής της δυσκολίας, πολλές χαλαρώσεις τους έχουν εισαχθεί τα τελευταία χρόνια. Επικεντρωνόμαστε στην, κατά πολλούς, πιο σημαντική: το κριτήριο EFX. Ως τώρα, το EFX υπάρχει μόνο σε πολύ περιορισμένες καταστάσεις, όπως όταν υπάρχουν το πολύ τρεις παίκτες ή όταν όλοι οι παίκτες έχουν την ίδια συνάρτηση αποτίμησης. Ακόμη και για την προσεγγιστική εκδοχή, ο καλύτερος γνωστός λόγος ισούται με $\phi - 1 (\approx 0.618)$

Η δουλειά μας οργανώνεται σε τρεις άξονες. Πρώτον, κατασκευάζουμε ένα πλαίσιο προσέγγισης για αθροιστικές συναρτήσεις αποτίμησης το οποίο ελέγχει την αλληλεπίδραση μεταξύ της ισχύος μια συνθήκης και της ποιότητας της προσέγγισης. Το κύριο αποτέλεσμά μας εδώ είναι λόγος προσέγγισης $2/3$ όταν οι παίκτες συμφωνούν στην κατάταξη των πρώτων αντικειμένων. Δεύτερον, προτείνουμε μια νέα μέθοδο για παρόμοιες κατατάξεις των αγαθών την οποία ονομάζουμε κατατάξεις σε στρώματα. Δείχνουμε ότι το EFX υπάρχει όταν τα στρώματα έχουν μέγεθος μέχρι 3, ακόμη και για πιο γενικές συναρτήσεις αποτίμησης από τις αθροιστικές. Τρίτον, εφαρμόζουμε τις νέες τεχνικές για να λάβουμε εναλλακτικές απλούστερες αποδείξεις για ορισμένα από τα υπάρχοντα αποτελέσματα. Ολοκληρώνουμε την εργασία με μία εμπειρική ανάλυση με πραγματικά δεδομένα, που μας δόθηκαν από την δημοφιλή ιστοσελίδα δίκαιου διαμοιρασμού Spliddit.

Λέξεις Κλειδιά

Αλγοριθμική θεωρία παιγνίων, ανάθεση πόρων, δίκαια διαμέριση, αδιαίρετα αγαθά, συναρτήσεις αποτίμησης.

Abstract

In this thesis, we study discrete fair division; that is allocating indivisible goods to agents in a fair manner. The problem is motivated by a wide range of applications, from distributing toys to kids to splitting an inheritance, where no sharing of items or monetary compensations are allowed. Since fairness is a hard concept to quantify, many notions have been defined throughout the years. Unfortunately, the presence of indivisible items renders them infeasible. As a countermeasure, a number of relaxations have been introduced more recently. We focused on the, arguably, most compelling one: finding allocations satisfying the EFX criterion. So far, EFX is guaranteed to exists only in very restricted settings; most notably when there are at most three agents or when they have identical valuations. Even for its approximation version, no progress has been made past $\phi - 1 (\approx 0.618)$.

Our work is along three axes. Firstly, we construct an approximation framework for additive valuations which controls tradeoffs between the strength of a condition and the quality of the approximation. Our main result here is a $2/3$ ratio assuming a common top ranking. Secondly, we propose a new method to capture similar rankings which we call tiered rankings. Within our model, we show that EFX exists when the size of the tier is at most 3, even for a broader than the additive class of valuation functions. Finally, we apply our new techniques to produce alternative simpler proofs for some existing results. We conclude the thesis with some real world data experiments, based on data obtained from the popular fair division website Spliddit.

Keywords

Algorithmic game theory, resource allocation, fair division, indivisible items, valuation functions.

Contents

List of Figures

Κεφάλαιο 1

Εκτεταμένη Ελληνική Περίληψη

Ξεκινούμε την παρούσα διπλωματική εργασία συνοψίζοντας στα ελληνικά το πρόβλημα του διακριτού δίκαιου διαμερισμού και τη συνεισφορά μας σε αυτό. Στο παρόν κεφάλαιο παρουσιάζονται οι βασικοί ορισμοί και τα κυριότερα αποτελέσματά μας, χωρίς αποδείξεις και τεχνικές λεπτομέρειες.

1.1 Το πρόβλημα

Για μια ομαλή εισαγωγή, ας θεωρήσουμε ένα πρόσφατο ζήτημα που δημιούργησε, μεταξύ άλλων, η πανδημία του Covid19: την έλλειψη εξοπλισμού Μονάδων Εντατικής Θεραπείας. Έστω, λοιπόν, ότι ο αναγνωστής είναι υπεύθυνος για την στελέχωση των νοσοκομείων με μηχανήματα οξυγόνου. Κάθε μηχάνημα πρέπει να σταλεί σε ένα νοσοκομείο, δηλαδή δεν είναι εφικτό να υπάρχει κάποιου είδους μοιρασιά αναμέσα σε περισσότερα νοσοκομεία. Φυσικά, δεν διαθέτουμε χρήματα για την αγορά επιπρόσθετων μηχανημάτων οπότε πρέπει τα υπάρχοντα να μοιραστούν με όσο πιο δίκαιο τρόπο για την κάλυψη των αναγκών κάθε περιοχής.

Μεταφράζοντας το παραπάνω πρόβλημα σε μαθηματική γλώσσα καταλήγουμε στον ακόλουθο ορισμό:

Ορισμός 1. Έστω ένα σύνολο παικτών $\mathcal{N} = \{1, \ldots, n\}$ με συναρτήσεις αποτίμησης $v_1, \ldots, v_n$ αντίστοιχα και ένα σύνολο αδιαίρετων αγαθών $\mathcal{M} = \{1, \ldots, m\}$. Μια ανάθεση $\mathcal{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ είναι οποιαδήποτε διαμέριση του $\mathcal{M}$ σε n υποσύνολα. Μια ανάθεση καλείται δίκαια αν ένα κριτήριο δικαιοσύνης ικανοποιείται για κάθε παίκτη.

Σε αυτό το σημείο, θα ήταν χρήσιμο να εξηγήσουμε την σημασία των συναρτήσεων αποτίμησης. Δεδομένου ότι κάθε άτομο είναι διαφορετικό, είτε μιλάμε για τον πραγματικό κόσμο είτε για το πρόβλημά μας, είναι επόμενο να ικανοποιείται σε διαφορετικό βαθμό από διαφορετικά αντικείμενα, π.χ., κάποιος μπορεί να προτιμάει ένα βιβλίο ενώ κάποιος άλλος μια ταινία. Επιστρέφοντας στο παράδειγμα του κορωνοϊού, το νοσοκομείο ενός μικρού χωριού θα είναι πλήρως ευχαριστημένο με λίγα μηχανήματα (και ίσως να μην έχει χώρο για την φύλαξη περισσοτέρων) ενώ ένα νοσοκομείο σε μεγάλο αστικό κέντρο θα ικανοποιηθεί πολύ λιγότερο με τον ίδιο αριθμό. Εν προκειμένω, αυτή η πληροφορία μας δίνεται μέσω των συναρτήσεων αποτίμησης.

Ορισμός 2. Για το πρόβλημα της δίκαιης διαμέρισης αδιαίρετων αγαθών μια συνάρτηση αποτίμησης είναι οποιαδήποτε αύξουσα συνάρτηση με πεδίο ορισμού το δυναμοσύνολο του $\mathcal{M}$ και πεδίο τιμών το $\mathbb{R}_{\geq 0}$. Επιπρόσθετα, $v(\emptyset) = 0$.

Να αναφέρουμε ότι οι συναρτήσεις αποτίμησης θεωρούνται μέρος της εισόδου και δεν μας απασχολεί πόσο εύκολο ή δύσκολο μπορεί να είναι στην πραγματικότητα να τις πληροφορηθούμε. Από δω και στο εξής, όταν μιλάμε για ένα σύνολο παικτών σιωπηλά θα συμπεριλαμβάνουμε και τις συναρτήσεις τους. Επίπλεον, πρέπει να μιλήσουμε και για την πιο συχνή κατηγορία τέτοιων συναρτήσεων, τις αθροιστικές. Πολύ απλά, αυτές εκφράζουν το πιο καθημερινό παράδειγμα όπου η αξία ενός συνόλου αγαθών είναι το άθροισμα της αξίας του κάθε αγαθού μεμονωμένα.

1.2 Περί δικαιοσύνης

Ενδεχομένως ο αναγνώστης να έμεινε με μια απορία από τον Ορισμό 1: τι εννοούμε όταν λέμε 'κριτήριο δικαιοσύνης'· Βασικά, σε φιλοσοφικό επίπεδο ακόμη αξίζει να αναρωτηθεί κανείς τι σημαίνει δικαιοσύνη. Εμείς θα αρκεστούμε στην παρουσίαση κάποιων προσπαθειών να οριστεί ποσοτικά η έννοια.

Ορισμός 3 (Ισονομία)**.** Μια ανάθεση $\mathcal{A}$ λέγεται ισόνομη αν $v_i(\mathcal{A}_i) = v_j(\mathcal{A}_j)$ για κάθε ζεύγος παικτών i, j.

Ορισμός 4 (Αναλογικότητα)**.** Μια ανάθεση $\mathcal{A}$ λέγεται αναλογική αν $v_i(\mathcal{A}_i) \geq \dfrac{v_i(\mathcal{M})}{n}$ για κάθε παίκτη i.

Ορισμός 5 (Ελεύθερη ζήλιας)**.** Μια ανάθεση $\mathcal{A}$ λέγεται ελεύθερη ζήλιας αν $v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j)$ για κάθε ζεύγος παικτών i, j.

Στον τελευταίο ορισμό αποδώσαμε τον αγγλικό όρο envy free που για συντομία συμπτίσσεται ως EF. Δεδομένου ότι η ελληνική απόδοση δεν είναι ιδιαίτερα εύηχη, θα κρατήσουμε την αγγλική συντομογραφία EF.

Και οι τρεις έννοιες αναπτύχθηκαν αρχικά για το διαμοιρασμό διαιρετών αγαθών. Σύντομα διαπιστώθηκε ότι στην περίπτωση των αδιαίρετων είναι μη επιτεύξιμες.

Παράδειγμα 1. Έστω ότι το $\mathcal{M}$ περιέχει m αντίγραφα του ίδιου αγαθού g. Αν το m είναι πολλαπλάσιο του n τότε η ανάθεση m/n αντιγράφων σε κάθε παίκτη πληροί και τις τρεις έννοιες δικαιοσύνης. Ωστόσο, σε κάθε άλλη περίπτωση καμία ανάθεση δεν ικανοποιεί καμία από τις τρεις έννοιες.

Για να ξεπεραστεί το παραπάνω εμπόδιο χρειάζεται να χαλαρώσουμε κάπως την έννοια της δικαιοσύνης που επιδιώκουμε. Προσπάθειες έγιναν για την έννοια της αναλογικότητας και την έλλειψη ζήλιας (EF). Επικεντρωνόμαστε στη δεύτερη κατεύθυνση. Μια πρώτη σκέψη για την υπερπήδηση του εμποδίου ήταν η εξάλειψη της ζήλιας έπειτα από μια θεωρητική αφαίρεση ενός αγαθού.

Ορισμός 6 (EF1)**.** Μια ανάθεση $\mathcal{A}$ λέγεται EF1 αν $\exists g \in \mathcal{A}_j : \ v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j \setminus g)$ για κάθε ζεύγος παικτών i, j.

Παράδειγμα 2. Έστω ένα στιγμιότυπο του προβλήματος με 2 όμοιους παίκτες και 3 αγαθά a, b, c με $v(a) = 10, \ v(b) = 5$ και $v(c) = 4$. Τότε, η ανάθεση του a στον έναν παίκτη και των b, c στον άλλο είναι EF1. Δυστυχώς, το ίδιο ισχύει και για την ανάθεση όπου τα a, c δίνονται στον ίδιο παίκτη.

Βλέπουμε, δηλαδή, ότι η χαλαρωμένη έννοια είναι υπερβολικά ασθενής, με αποτελέσμα να επιτρέπει διαισθητικά άδικες αναθέσεις. Παρ' όλ' αυτά, μπορούμε να την αυστηροποιήσουμε.

Ορισμός 7 (EFX)**.** Μια ανάθεση $\mathcal{A}$ λέγεται α-EFX αν $\forall g \in \mathcal{A}_j : \ v_i(\mathcal{A}_i) \geq \alpha \cdot v_i(\mathcal{A}_j \setminus g)$ για κάθε ζεύγος παικτών i, j.

Στον προηγούμενο ορισμό, απλώς, δώσαμε δύο έννοιες μαζί, αυτή του κριτηρίου EFX για $\alpha = 1$ και της προσεγγιστικής εκδοχής για τιμές μικρότερες της μονάδας. Το κριτήριο EFX αποτελεί τον κύριο στόχο στην ερευνητική περιοχή όπου εντάσσεται η διπλωματική και όλες μας οι συνεισφορές περιστρέφονται γύρω από αυτό.

1.3 Γνωστοί αλγόριθμοι και αποτελέσματα

Ίσως ο γνωστότερος αλγόριθμος στην ερευνητική περιοχή που εξετάζουμε είναι ο αλγόριθμος Αποκλεισμού Κύκλων Ζήλιας (Envy Cycle Elimination).

Αλγόριθμος 1 Αποκλεισμού Κύκλων Ζήλιας

Θέσε $\mathcal{A}_i = \emptyset$ για κάθε παίκτη i
while $\exists$ κάποιο μη ανατεθέν αγαθό g **do**
  while $\exists$ κάποια πηγή s **do**
    Θέσε $\mathcal{A}_s = \mathcal{A}_s \cup g$
  end while
  Κατέστησε τον γράφο ακυκικλό μετακινώντας τα μερίδια των παικτών πάνω σε έναν κύκλο ζήλιας όπως ορίζει ο κύκλος
end while
return $\mathcal{A} = (\mathcal{A}_1, \ldots, \mathcal{A}_n)$

Τα κυριότερα αποτέσματα στο κυνήγι EFX αναθέσεων είναι:

- Ύπαρξη για n παίκτες με μία από δύο γενικές συναρτήσεις αποτίμησης
- Ύπαρξη για n παίκτες με κοινή κατάταξη των αγαθών και αθροιστικές συναρτήσεις
- Ύπαρξη για 3 παίκτες με αθροιστικές συναρτήσεις
- Ύπαρξη για $m = n + 3$ αγαθά
- Πρόσεγγιση $\phi - 1 \approx 0.618$ για n παίκτες με αθροιστικές συναρτήσεις

1.4 Συνεισφορά

Για την καλύτερη οργάνωση της συνεισφοράς μας, την χωρίζουμε με βάση την οικογένεια συναρτήσεων αποτίμησης.

Αθροιστικές συναρτήσεις

Αλγόριθμος 2 Γενικό Προσεγγιστικό Πλαίσιο

Αρχικοποίησε με την κενή ανάθεση
Υπολόγισε μία μερική α-EFX ανάθεση S διατηρώντας την ιδιότητα

$$v_i(S_i) \geq \beta \cdot v_i(h) \text{ για κάθε } h \in M \setminus S$$

Συνέχισε με τον αλγόριθμο 1 εώς ότου εξαντληθούν τα αντικείμενα

Θεώρημα 1. Ο αλγόριθμος 2 υπολογίζει μια $\min\left(\alpha, \frac{\beta}{\beta+1}\right)$-EFX ανάθεση.

Το προσεγγιστικό πλαίσιο μας δίνει τη δυνατότητα να ρυθμίζουμε το λόγο προσέγγισης ανάλογα με τον βαθμό που ισχύει η συνθήκη ύπαρξης. Για παράδειγμα, αν οι παίκτες δεν έχουν κοινή κατάταξη για όλα τα αγαθά αλλά έχουν για τα πρώτα $4n$ το πλαίσιο μας εγγυάται λόγο προσέγγισης $4/5$. Η ελάχιστη συνθήκη θα ήταν αν είχαν κοινή κατάταξη στην πρώτη n-αδα με λόγο $1/2$. Ωστόσο, παρατηρούμε ότι για την επίτευξη αυτού του λόγου δεν χρειάζονται υποθέσεις.

Λήμμα 1. Ξεκινώντας από την μερική ανάθεση όπου κάθε παίκτης διαλέγει με την σειρά το αγαπήμενό του διαθέσιμο αγαθό, ο αλγόριθμος 2 υπολογίζει μία $1/2$-EFX ανάθεση.

Συνεπώς, ήταν λογικό με την επιπρόσθετη συνθήκη για την πρώτη n-αδα να μπορούμε να επιτύχουμε ακόμα καλύτερο λόγο προσέγγισης.

Αλγόριθμος 3 $2/3$ EFX για κοινή κατάταξη των n πρώτων αγαθών

Επίλεξε μία αυθαίρετη σειρά των παικτών
Δώσε σε κάθε παίκτης με τη σειρά το αγαπημένο του διαθέσιμο αγαθό
for i in range(N,1) **do**
  Όρισε $\mathcal{A}_i = \mathcal{A}_i \cup \arg\max\ v_i(g)$
  Επίλεξε τον παίκτη $j = \arg\max_{k \in [i]}\ v_i(\mathcal{A}_k)$
  if $i \neq j$ **then**
    $\mathcal{A}_i = \mathcal{A}_j$
    $\mathcal{A}_j = \arg\max\ v_j(m)$
    Μετακίνησε τους παίκτες $i-1$ ως $j+1$ κατά μία θέση
    Μετακίνησε τον j στην τρέχουσα πρωτη θέση
  end if
end for
Συνέχισε με τον αλγόριθμο 1

Θεώρημα 2. Ο αλγόριθμος 3 υπολογίζει μία 2/3-EFX ανάθεση όταν οι παίκτες συμφωνούν στην κατάταξη των n πρώτων αγαθών.

Απλοποιήσιμες Συναρτήσεις

Ορισμός 8. Μία συνάρτηση αποτιμήσης v ονομάζεται απλοποιήσιμη εάν για οποιαδήποτε πακέτα αγαθών $S, T \subset M$ και οποιοδήποτε αγαθό $g \in M \setminus (S \cup T)$ ισχύει ότι

$$v(S \cup g) > v(T \cup g) \implies v(S) > v(T)$$

Βλέπουμε ότι στον ορισμό αυτής της κλάσης εμφανίζονται τα σύνολα S και $S \cup g$. Επί της ουσίας, το ίδιο συμβαίνει και στον ορισμό του EFX όπου εμφανίζονται τα $\mathcal{A}_i$ και $\mathcal{A}_i \setminus g$. Συνεπώς, φαίνεται να είναι βολική για την μελέτη του προβλήματος. Επιπρόσθετα, εύκολα διαπιστώνουμε ότι οι αθροιστικές συναρτήσεις είναι απλοποιήσιμες άρα η κλάση αποτελεί ένα ενδιάμεσο βήμα πριν τις γενικές συναρτήσεις αποτίμησης.

Σε αυτή την διπλωματική σημειώσαμε πρόοδο στην μελέτη του προβλήματος για τις απλοποιήσιμες συναρτήσεις, ξεκινώντας από τον παρακάτω ορισμό.

Ορισμός 9. Μία μερική κατάταξη $\mathcal{T}$ είναι μια κοινή κατάταξη *σε στρώματα* αν υπάρχει μία διατεταγμένη διαμέριση όλων των αντικειμένων $M = (M_1, M_2, \ldots, M_l)$ τέτοια ώστε:

$$\forall g \in M_i, \forall h \in M_{j>i} : \ v(g) \geq v(h)$$

για όλους τους παίκτες. Επιπλέον, ορίζουμε ως μέγεθος του $\mathcal{T}$ το μέγεθος του μεγαλύτερου M_i.

Παρατηρείστε ότι μία κατάταξη σε στρώματα μεγέθους 1 είναι απλώς η κοινή κατάταξη. Συνεπώς, ο ορισμός μας επεκτείνει ένα γνωστό πλαίσιο. Αξιοποιώντας τον όρισμο δείξαμε το εξής αποτέλεσμα.

Θεώρημα 3. Αν όλοι οι παίκτες έχουν απλοποιήσιμες συναρτήσεις αποτίμησης και κοινή κατάταξη των αγαθών σε στρώματα μεγέθους το πολύ 3, τότε υπάρχει πάντοτε μια EFX ανάθεση.

Για την απόδειξη του παραπάνω θεωρήματος χρησιμοποιήσαμε και ένα λήμμα το οποίο και θα διατυπώσουμε εδώ καθώς μπορεί να ειναι ανεξάρτητου ενδιαφέροντος.

Λήμμα 2 (Πρόσθεση ανισοτήτων)**.** Έστω S, T, Q και R σύνολα τέτοια ώστε $S \cap Q = \emptyset$ και $T \cap R = \emptyset$. Τότε

$$\left.\begin{aligned} v(S) \geq v(T) \\ v(Q) \geq v(R) \end{aligned}\right\} \implies v(S \cup Q) \geq v(T \cup R)$$

1.5 Επίλογος

Κλείνοντας αυτήν την ελληνική περίληψη, πρέπει να αναφέρουμε ότι προσπαθήσαμε να ελέγξουμε εμπειρικά τα αποτελέσματά μας, μέσω δοκιμών σε πραγματικά δεδομένα τα οποία λάβαμε από την ιστοσελίδα `Spliddit.org`. Δυστυχώς, το συμπέρασμα ήταν ότι η πράξη απέχει, ακόμη, από την θεωρία. Για τα αναλυτικά γραφήματα παραπέμπουμε τον αναγνωστή στο κεφάλαιο 6 του αγγλικού κειμένου. Τέλος, αφήνουμε ως ανοικτές ερωτήσεις για μελλοντική εργασία τόσο την ύπαρξη 2/3-EFX αναθέσεων άνευ συνθηκών όσο και την επέκταση στο μέγεθος των στρωμάτων.

Chapter 2

Introduction

2.1 Motivation

Dividing resources among people in a manner that satisfies everyone is not exactly a new problem. In the Bible, upon arriving in Canaan, Abraham and Lot set to fairly divide the land between them. In more recent examples, the reader of this thesis may have tried to divide a collection of NFTs among their friends. Of course, all of them are treated as equals meaning that no one should feel dissatisfied by their share.

The first study of fair division from a mathematical viewpoint dates back to Steinhaus in 1948, [Ste48]. The problem there was illustrated via cutting a cake; making the phrase a synonym of fair division. Imagine we have a dual flavoured cake: vanilla and chocolate, and it must be split into two pieces, one for Alice and one for Bob. However, there is a catch: Alice prefers vanilla and Bob prefers chocolate. Therefore, weighting the cake and giving half to each child is not enough. Alice may receive the piece with more chocolate and Bob the one with more vanilla. The proposed solution here is to have Alice cut the cake into pieces with equal amounts of vanilla and let Bob choose the one he likes the most.

At this point, the reader may wonder what happens if a third child comes to the party. Is it still possible to partition the cake in a fair manner? The short answer would be "yes" but the follow up question would be "how"? The point of this little discussion is to emphasize the difference between existence (the answer is "yes") with computation ("how"). In the age of personal computers, most mathematical problems have followed this very path: from theory to algorithms. In the first example of Alice and Bob, the cut-and-choose method was the algorithm, albeit a simple one. Our work will follow this path as well: every existential result will be paired with an algorithm computing the promised allocation.

Next, we have to establish some division rules. Imagine that we have a land like Abraham and Lot. Can we make a contract and give, for example, 55% to Alice? That would result in a kind of co-ownership. Or is there any spare cash to give to Bob to compensate for his smaller share? In many occasions, the answer would be affirmative in both scenarios. Indeed, a lot of work has been carried over the years in various such settings, see, for instance, [ADG91] for settings with money and [Mou04] for joint ownership. In general, in the previous century the problem was mainly studied for divisible goods, e.g., [RW98] and [Mou04], with indivisible ones only studied when some divisible resources were also at hand.

However, in the last decade or so, a new paradigm has emerged causing a surge of research in the area. To illustrate it, think about a recent issue caused by Covid19: a shortage on ventilators. In this scenario, we are in charge to distribute a number of the life saving machines to a number of hospitals. Clearly, we cannot send 55% of a ventilator to a hospital. Neither can we have the two hospitals sharing a machine with time shifts; during the transport from one hospital to another valuable time would be lost. Similarly, we cannot compensate a hospital that does not have enough ventilators with money or otherwise. One way or another, we must choose which hospital to send the equipment to.

To conclude this short introduction, there are many fair division types of problems, depending on whether we can compensate agents with currency and on whether or not the goods can be split/shared. Inspired from problems like the one above, we will focus only on instances where each item must be given as a whole to a single agent.

2.2 Contribution

So far, we have used the term "fair" somewhat loosely, invoking the reader's inner sense of justice. For a mathematical study of the problem, this will not suffice. Thus, we will present the most prominent fairness notions with their pros and cons in order to establish which one is most suitable for our problem. As we will argue in the sequel, this will be the notion of *envy freeness up to any good*, [CKM+16], or EFX for short.

Having decided upon the fairness notion, we then proceed with the literature in two ways. Firstly, we present the main tools developed: *Greedy Round Robin* and *Envy Cycle Elimination*, [LMMS04]. Secondly, we present an important part of the current literature, mainly centered around the existence of EFX allocations for *identical* valuations, [PR18], and EFX with *bounded charity*, [CKMS21]. Given that this a very active line of research, a complete survey is out of the scope of this thesis.

Apart from bibliographical insights, we also provide some new theoretical results.

- In the study of additive valuations, we show a simpler way to match the ratio of 1/2 in the approximate version of EFX. Based on that result we show how simple conditions can allow us to improve the ration even further. The culmination of our work leads to a general approximation framework that can leverage conditions to get better approximations. The highlight of the framework is a 2/3-EFX allocation when the agents have a same ranking for their (few) top items.

- Since ordinal information seems fruitful in improving the approximation ratio, we then explore what can we say about similar rankings. We introduce the term *tiered* *ranking* to denote a partial ranking with some missing information, allowing the agents to deviate from the identical setting. We then prove that an EFX allocation exists when the size of the tier is at most 3; a result that holds for the broader class of cancelable valuations.

– Finally, for general valuations we present a new allocation rule which allows Envy Cycle Elimination to compute an EFX allocation not only when agents have identical valuations but also when they have identical rankings over all item subsets. The same rule is also used to obtain a non cut-and-choose algorithm for two players and an EFX allocation when the number of items is at most two greater than the number of agents.

2.3 Document outline

The structure of the thesis is:

Chapter 3: A formal definition of the problem followed by a discussion about the different fairness notions and their relaxations

Chapter 4: An overview of the two main algorithms in the area and and a summary of the currently known results about EFX allocations

Chapter 5: Our contributions

Chapter 6: Experimentation with real world data

Chapter 7: Conclusion with some directions for future work

Chapter 3

Fundamentals of Discrete Fair Division

Discrete Fair Division lies within the broader field of Algorithmic Game Theory. As such, it concerns scenarios where a decision needs to be made for *players* or *agents* with different goals and objectives. In the language of Game Theory, an agent i is described by her *valuation function* v_i, which reflects her opinion about a possible output of some algorithm; in our case, the output is, of course, some bundles of items.

3.1 Discrete Fair Division

There are two major Fair Division settings: the *continuous* and the *discrete* one. Before proceeding any further, we feel the need to explain the difference that we already touched in the Introduction to a greater extent. To that end, consider the most famous example of the continuous setting: cutting a cake. The cake may consist of many layers - chocolate, caramel etc. - and each agent prefers different flavours to different extents (a first example of valuation functions). Intuitively, we can cut the cake to as many and as tiny pieces as needed to accommodate all the agents. Now, replace the cake with some jewels. Unfortunately, cutting one jewel deprives it of value. The fact that cutting one item may destroy it or, in other words, that the item is *indivisible*, is the distinction between the two settings.

While the continuous case is not as trivial as our previous example made it look, it is evident that Discrete Fair Division is more challenging, and thus unexplored. On the contrary, there is a vast literature on cake cutting and continuous Fair Division. The interested reader may start with [Pro15] or [BT96].

On a different note, observe that in both the examples mentioned above, regarding the cake and the jewels, the goods under consideration are desired by the agents. This is not always the case in fair division problems. For example, if each agent is a CPU and the items to be allocated are processes, similarly to a job scheduling scenario, the agents are not pleased when receiving a greater load of work. That means the items can be divided into *goods* and *chores*, as named in the literature. In this thesis we consider only the division of goods, and the terms "items" and "goods" are going to be used interchangeably. For more about *chores* we refer the reader to [ACIW22] and references therein.

At this point we can make a first attempt to formulate the problem.

Definition 3.1.1 (Discrete Fair Division)**.** Let $\mathcal{N} = \{1, \ldots, n\}$ be a set of agents and $\mathcal{M}$ be a set of indivisible goods. An *allocation* $\mathcal{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ is any valid partition of $\mathcal{M}$; that is $\mathcal{A}_i \cup \mathcal{A}_j = \emptyset$ for every pair of agents i, j and $\bigcup_{i \in \mathcal{N}} \mathcal{A}_i = \mathcal{M}$. An allocation is *fair* if a fairness criterion is satisfied for every agent.

That begs the question: *what is fairness*?

We will (try to) answer the question, but firstly we must note something about Definition 3.1.1: it informs us that we are dealing with sets and subsets. Therefore, the domain of every valuation function v should be the powerset of $\mathcal{M}$, or $v : 2^{\mathcal{M}} \to \mathbb{R}$. However, not all of these functions make sense in the context we are studying. If we loosely translate v to the "happiness" of an agent we would expect it to improve every time the agent receives more items. Similarly, we could wonder what is the appropriate value of "happiness" when some agent i does not participate, i.e., $\mathcal{A}_i = \emptyset$. Addressing those issues leads to the following definition.

Definition 3.1.2 (General valuation function)**.** In the context of Fair Division with goods, a *general* valuation function $v : 2^{\mathcal{M}} \to \mathbb{R}_{\geq 0}$ must obey the following two assumptions:

i) *normalization*, i.e., $v(\emptyset) = 0$

ii) *monotonicity*, i.e., $S \subseteq T \implies v(S) \leq v(T)$

In the vast majority of the literature, the valuation functions are actually restricted to a more "everyday" class:

Definition 3.1.3 (Additive valuation function)**.** A valuation function $v : 2^m \to \mathbb{R}_{\geq 0}$ is additive if $v(S) = \sum\limits_{g \in S} v(g)$, $\forall S \subseteq M$.

3.2 Fairness notions

At this point we can return to our question about fairness. Leaving aside our theories for a moment, if we were to go outside and ask people "what is fairness" the most probable answer would be something along the lines of "everyone should receive the same".

Definition 3.2.1 (Equitability)**.** An allocation $\mathcal{A}$ is equitable if $v_i(\mathcal{A}_i) = v_j(\mathcal{A}_j)$ for any pair of agents i, j.

Note the different subscripts of the valuation functions: agent i may not necessarily value $\mathcal{A}_i$ and $\mathcal{A}_j$ equally but if we ask the agents how much is their bundle worth to them, we get the same answer from everyone. Unfortunately, this is an extremely stringent requirement. To verify it, just consider an instance with 2 agents and 2 items where the agents do not prefer the same item. The next fairness notion can be seen as asking only one agent at a time. Assuming that Alice knows that all the goods combined cost 100$ and there are ten agents in total, what would be a fair value for her bundle?

Definition 3.2.2 (Proportionality)**.** An allocation $\mathcal{A}$ is proportional if $v_i(\mathcal{A}_i) \geq \frac{v_i(M)}{n}$.

For this notion, note the use of inequality instead of strict equality. Continuing with our toy example, Alice would be expecting a 10$ bundle. Give her a 9$ one and she will most definitely complain but no one will complain were they to receive more. One the other hand, what if Alice received a bundle she values at 11$ but Bob got one at 12$?

Definition 3.2.3 (Envy-freeness)**.** Let $\mathcal{A}$ be an allocation and i, j a pair of agent such as $v_i(\mathcal{A}_i) < v_i(\mathcal{A}_j)$. Then we say that agent i *envies* agent j and denote this by $i \rightarrow j$. If no such pair exists, $\mathcal{A}$ is *envy-free.*

3.3 The need for relaxations

With the fairness notions in place, our next goal should have been designing algorithms to compute them. Unfortunately, we will show that these notions are too strong to ask for. Consider the most minimal example possible: two agents and an odd number of identical goods g. Equitability cannot be achieved since one agent will value a bundle at least $v(g)$[1] more . Neither can proportionality since the copies of g are an odd number and, finally, whichever side will be receiving less items will be envying the other, thus ruling out envy-freeness.

So we have established the need for some relaxed notions of fairness, justifying this section's title. Still, one may wonder if it is possible to identify restricted settings where the strong notions can be satisfied and if it is fruitful to search for allocations in those settings. The answer is a definite "No".

Proposition 1 ([LMMS04])**.** *Computing an allocation satisfying any of the three notions of fairness is computationally intractable.*

Proof sketch. Consider a setting with 2 agents and identical additive valuation functions. Then, the problem of satisfying any notion reduces to *partitioning* an array of integers into two parts of equal sum; thus it cannot be solved efficiently, unless P=NP. □

To the best of our knowledge, no attempts to relax equitability have been proposed in the literature. This is probably because envy-freeness is already a relaxed and more

[1] Note that we abused the notation, writing $v(g)$ instead of $v(\{g\})$; a practice we will follow throughout this thesis.

natural version[2]. Therefore, we shift our focus to relaxing the notion of envy-freeness. We saw that problems arise even in the simplest setting with one good and two agents. As a result, any proposed relaxation should start by addressing this scenario. The first try was by Budish, [Bud11].

Definition 3.3.1 (EF1)**.** An allocation $\mathcal{A}$ is *envy-free up to **1** good* (EF1) if for every pair of agents i, j it holds that

$$\exists g \in \mathcal{A}_j : \ v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j \setminus g)$$

The introduction of EF1 trivially solves the problem with 2 agents and odd copies of g since $\mathcal{A}_i \setminus g = \mathcal{A}_j$. Moreover, as we will see later in paragraph 4.1.2, EF1 allocations always exist and can be computed efficiently. So, are we done? Well, not quite. Let us first examine the following example.

Example 3.1 (The unfairness of EF1)**.** Assuming that we have two agents with identical additive valuations v and 3 goods: $v(g_1) = 10$, $v(g_2) = 5$ and $v(g_3) = 4$. Then, $\{g_1, \{g_2, g_3\}\}$ is an EF1 allocation since agent 2 will not envy $\mathcal{A}_1 = g_1$ after the removal of g_1. However, there is one more EF1 allocation: $\{g_2, \{g_1, g_3\}\}$ since $v_1(g_2) > v_1(g_3) = v_1(\mathcal{A}_2 \setminus g_1)$. △

Clearly, anyone diving that set of 3 goods, and without any fair division knowledge, would pick the first allocation and would not even consider the second. Which means that our relaxed notion of fairness is actually too relaxed. Caragiannis et al., [CKM+16], proposed a stricter version of EF1.[3]

Definition 3.3.2 (EFX)**.** An allocation $\mathcal{A}$ is *envy-free up to **any** good* (EFX) if for every pair of agents i, j it holds that

$$\forall g \in \mathcal{A}_j : \ v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j \setminus g)$$

A subtle yet crucial change from the definition of EF1: by strengthening the requirement from dropping some good to dropping any good the "unfair" allocation of example 3.1 is ruled out. Similarly to EF1 and EFX, one can define the equivalent relaxations of proportionality, namely Prop1 and PropX.

The next relaxed fairness notion is based on a different train of thought. Imagine that instead of dividing the goods ourselves we give them to Alice and ask her to help us. At first, that seems flawed since Alice may try to keep everything for herself. But there is a catch: once she divides the items into bundles, she will be the last one to pick. To counter that disadvantage, Alice must try to make even her least favorite bundle valuable; thus making the allocation, in a way, fairer. The following definition is due to Budish, [Bud11].

[2]Consider a 2-agent setting with the goods a, b where agent 1 prefers a and 2 prefers b. Allocating the goods based on agents' preference gives us an envy-free but not equitable allocation. Should the latter be an issue?

[3]The notion was already introduced by Gourvès et al.,[GMT14], under the name *near envy-freeness*.

Definition 3.3.3 (MMS)**.** Let $\mathcal{A}(\mathcal{N},\mathcal{M})$ be the set of all possible allocations of the goods in $\mathcal{M}$ to the agents in $\mathcal{N}$. The *maximin share* $\mu(\mathcal{N},\mathcal{M})$ of agent i is defined as

$$\mu_i(\mathcal{N},\mathcal{M}) = \max_{\mathcal{A}\in\mathcal{A}(\mathcal{N},\mathcal{M})} \min_{\mathcal{A}_i\in\mathcal{A}} v_i(\mathcal{A}_i)$$

An allocation $\mathcal{A}$ is *maximin share fair* (MMS) if for every agent $i \in \mathcal{N}$ it holds that

$$v_i(\mathcal{A}_i) \geq \mu_i(\mathcal{N},\mathcal{M})$$

When it is clear from context, we will drop the $(\mathcal{N},\mathcal{M})$ part of the notation and simply refer to the maximin share of i as μ_i. MMS solves the problem of dividing $2k+1$ identical items in a different manner: since an agent must create two bundles, one will have k copies and the other $k+1$ and both agents will be satisfied with either. Unfortunately, even if MMS is a relaxed notion itself, it is still too strict to always exist.

Example 3.2 (Example 7 of [BL16])**.** Consider the following instance with 4 goods $\mathcal{M} = \{a,b,c,d\}$ and 2 agents with valuation functions:

$$v_1(S) = \begin{cases} 1, \text{ if } S = \{a,b\},\{c,d\} \text{ or } |S| \geq 3 \\ 0, \text{ otherwise} \end{cases}$$
$$v_2(S) = \begin{cases} 1, \text{ if } S = \{a,c\},\{b,d\} \text{ or } |S| \geq 3 \\ 0, \text{ otherwise} \end{cases}$$

It is easy to see that agent 1 would partition the goods into $\{a,b\}$ and $\{c,d\}$, yielding $\mu_1 = 1$. Likewise, $\mu_2 = 1$ due to the partition $\{a,c\}/\{b,d\}$. However, no allocation can give their maximin share to both agents simultaneously. △

As a consequence, variations of MMS have become a new research direction, as initiated in [CKM+16].

Definition 3.3.4 (Pairwise MMS)**.** An allocation $\mathcal{A}$ is *pairwise maximin share fair* (PMMS) if for every pair of agents $i,\ j \in \mathcal{N}$ it holds that

$$v_i(\mathcal{A}_i) \geq \max_{\mathcal{B}\in\mathcal{A}(\{i,j\},\mathcal{A}_i\cup\mathcal{A}_j)} \min(v_i(\mathcal{B}_i), v_i(\mathcal{B}_j))$$

In words, the pairwise maximin share of agent i is computed by merging her bundle with j's and then redistributing the cumulative share in an MMS fashion. Of course, one can go a step further and repeat the same process with any group of agents, [BBKN18], instead of just pairs.

Definition 3.3.5 (Groupwise MMS)**.** An allocation $\mathcal{A}$ is *groupwise maximin share fair* (GMMS) if for every group of agents $G \subseteq \mathcal{N}$ it holds that

$$v_i(\mathcal{A}_i) \geq \max_{\mathcal{B}\in\mathcal{A}(G,\cup\mathcal{A}_j)} \min_{\mathcal{B}} v_i(\mathcal{B}_j)$$

For other relevant fairness notions, we refer the reader to Chapter 5 of the recent survey by Amanatidis et al., [ABFRV22].

3.4 Approximate versions

As mentioned, we will solely focus on the pursuit of EFX. It is, after all, "fair division's most enigmatic question" - Procaccia, [Pro20]. As is common when dealing with Computer Science problems, figuring out an exact answer is difficult, if not outright impossible. Thus one seeks approximate answers.

Definition 3.4.1 (α-EFX)**.** An allocation $\mathcal{A}$ is α-EFX if for every pair of agents i, j it holds that

$$\forall g \in \mathcal{A}_j : \; v_i(\mathcal{A}_i) \geq \alpha \cdot v_i(\mathcal{A}_j \setminus g)$$

Obviously, setting $\alpha = 1$ retrieves definition 3.3.2. Our new objective is to compute allocations satisfying the property with α as closer to the unity as possible. Of course, one can respectively define α-EF1, α-MMS etc. For a detailed comparison of the approximate notions we refer the reader to [ABM18]. From there we will also borrow the final example of this chapter.

3.4.1 A thorough example

Example 3.3 (Example 1 of [ABM18])**.** Consider the instance with 3 agents, $M = \{a, b, c, d, e\}$ a set of 5 goods and valuations:

	a	b	c	d	e
Agent 1	3	1	1	1	4
Agent 2	4	3	3	1	4
Agent 3	3	2	1	3	4

We start by examining the maximin shares of the agents. It is easy to check that:

- $\mu_1 = 3$ due to the partition $\{a\}$ / $\{b, c, d\}$/ $\{e\}$
- $\mu_2 = 4$ due to the partition $\{a\}$ / $\{b, c\}$ / $\{d, e\}$
- $\mu_3 = 4$ due to the partition $\{a, b\}$ / $\{c, d\}$ / $\{e\}$

Of course, the MMS guaranteeing partitions are not unique. Moving on to allocations, the allocation $\mathcal{A} = (\{e\}, \{b, c\}, \{a, d\})$ is envy-free, and thus EF1 and EFX, and MMS, thus also PMMS and GMMS. On the other hand, the allocation $\mathcal{B} = (\{a\}, \{b, e\}, \{c, d\})$ is only EF1 and MMS. We will leave the details as an exercise to the reader and proceed with the calculation of α from definition 3.4.1.

- Agent 1

$$v_1(\mathcal{A}_1) = 3 \leq 4 = v_1(e) = v_1(\mathcal{A}_2 \setminus b) \implies \alpha \leq 3/4$$
$$v_1(\mathcal{A}_1) \geq v_1(\mathcal{A}_3) \implies \alpha \leq 1$$

- Agent 2

$$\begin{aligned} v_2(\mathcal{A}_2) \geq v_2(\mathcal{A}_1) &\implies \alpha \leq 1 \\ v_2(\mathcal{A}_2) \geq v_2(\mathcal{A}_3) &\implies \alpha \leq 1 \end{aligned}$$

- Agent 3

$$\begin{aligned} v_3(\mathcal{A}_3) \geq v_3(\mathcal{A}_1) &\implies \alpha \leq 1 \\ v_3(\mathcal{A}_3) = 4 \geq 4 = v_3(e) = v_3(\mathcal{A}_2 \setminus b) &\implies \alpha \leq 1 \end{aligned}$$

Combining the inequalities yields that $\mathcal{B}$ is a $\frac{3}{4}$-EFX allocation. $\triangle$

In the previous example, agent 1 is still envious of agent 3 even after the removal of item b. On the contrary, agent 3's envy towards 2 was eliminated. When pursuing exact EFX ($\alpha = 1$), it is useful to be able to separate the two cases. To that end, we say that some agent *strongly envies* another when the envy persists even after removing the least significant good.

Chapter 4

Literature Review

In this chapter, our aim is twofold. Firstly, we present the existing techniques based upon which we will build our contributions. Secondly, we discuss existing results to understand how ours are positioned within the current research agenda.

4.1 The main techniques

Despite many years of active research on the topic, there are only two main techniques used in pretty much every work, just with some minor tweaks every time. Those are the greedy round robin and the envy cycle elimination algorithms. Both (efficiently) compute EF1 allocations and are used as building blocks in the search for EFX.

4.1.1 Greedy Round Robin

The first such technique is a classic *round robin* algorithm. Although it appears in various works, e.g., [Mar17], it is not attributed to someone due to its simplicity. The algorithm is presented below.

Algorithm 1 Greedy Round Robin($\mathcal{N}, \mathcal{M}$)

1: Set $\mathcal{A}_i = \emptyset$ for every agent i
2: Fix some arbitrary agent ordering π
3: **while** $\exists$ some unallocated item **do**
4:     Let $i \in \mathcal{N}$ be the next agent according to π in a round robin fashion
5:     Let g be i's most preferred item among the currently unallocated items
6:     Set $\mathcal{A}_i = \mathcal{A}_i \cup g$
7: **end while**
8: **return** $\mathcal{A} = (\mathcal{A}_1, \ldots, \mathcal{A}_n)$

The term "greedy" is due to line 5: every agent picks greedily when it is her turn. Before proving the correctness of the algorithm we must note that it does not work for general valuations.

Proposition 2. *When the agents have additive valuations, Greedy Round Robin computes an EF1 allocation in polynomial time.*

Proof. Let i, j be any two agents such as i comes before j in π and let r be the number of rounds where i receive some good. Then we can write:

$$\mathcal{A}_i = (g_1, g_2, \ldots, g_{r-1}, g_r)$$
$$\mathcal{A}_j = (h_1, h_2, \ldots, h_{r-1}, h_r)$$

where h_r may not exist (if the algorithm run out of items between i's and j's turns), and we then treat it as a zero.

Now, $v_i(g_k) > v_i(h_k)$ for every k since agent i could have picked either item and by adding the inequalities $v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j)$; thus there is no envy from i's side. From j's side it may exist but it can be bounded based on the following observation: j prefers h_1 over g_2, h_2 over g_3 etc since now she was the one with the choice. Again by the adding the inequalities we receive $v_j(\mathcal{A}_j) \geq v_j(\mathcal{A}_i \setminus g_1)$ thus the allocation is EF1. The efficiency of the algorithm is trivial since we simply need m max operations. □

4.1.2 Envy Cycle Elimination

The second algorithm of this section is called Envy Cycle Elimination and is due to Lipton et al., [LMMS04]. It is based on a graph theoretic approach to the problem: consider a graph where each node represents an agent. Two nodes, or agents, i, j are connected with a directed edge $i \to j$ if and only if i envies j. This directed graph is called the *envy graph* of the allocation and is usually denoted by E_G. The algorithm works in incremental style, allocating one item at a time to some unenvied agent, and moves the bundles around when there is none.

Algorithm 2 Envy Cycle Elimination($\mathcal{N}, \mathcal{M}$)

1: Set $\mathcal{A}_i = \emptyset$ for every agent i
2: **while** $\exists$ some unallocated item g **do**
3:     **if** $\exists$ some source s **then**
4:         Set $\mathcal{A}_s = \mathcal{A}_s \cup g$
5:     **end if**
6:     Decycle the envy graph by repeatedly finding envy cycles and reallocating backword the bundles along the edges of each cycle
7: **end while**
8: **return** $\mathcal{A} = (\mathcal{A}_1, \ldots, \mathcal{A}_n)$

Let us explain line 6 (the decycling step) a bit more. Since the inner while loop broke we know that there are no sources in E_G. Or, in other words, every node has an incoming edge which mean that the graph has an *envy cycle* of the form agent $c_1 \to$ agent $c_2 \to \cdots \to$ agent c_l for some $l \geq 2$. In terms of envy, c_1 envies c_2, c_2 envies c_3 etc. Now, if we reallocate $\mathcal{A}_{c_2}$ to agent c_1, $\mathcal{A}_{c_3}$ to c_2 and so on until we allocate $\mathcal{A}_{c_1}$ to c_l, every agent on the cycle has improved their bundle while no agent outside it has gotten worse. Intuitively, that is a good measure of progress and it implies that the decycling process cannot go on forever.

Definition 4.1.1 (Pareto)**.** Let $\mathcal{A}$ and $\mathcal{B}$ be two possible allocations between the same group of agents $\mathcal{N}$. If $\forall i \in N: \; v_i(\mathcal{A}_i) \geq v_i(\mathcal{B}_i)$ with strict inequality for at least one agent we say that $\mathcal{A}$ *Pareto dominates* $\mathcal{B}$. Also, we say that $\mathcal{A}$ is *Pareto optimal* if no allocation dominates it.

We are ready to prove the following claim:

Proposition 3. *Envy Cycle Elimination computes an EF1 allocation in polynomial time.*

Proof. The valuation of any agent i is upper bounded by $v_i(\mathcal{M})$. Since eliminating the envy cycle in line 6 produces Pareto dominating allocations we are guaranteed to enter the inner loop eventually. Now, allocating any item to any agent also produces a new allocation that Pareto dominates the current one. Thus the algorithm must terminate after allocating all the items. Having proved the termination, showing the EF1 property is now trivial: $\mathcal{A}_s$ was not envied by any agent. Even if $\mathcal{A}_s \cup g$ is, removing g solves the issue. To complete the proof note that locating and removing a cycle costs $O(n^2)$, while at the same time reducing the number of edges in the graph. On the other hand, each newly allocated good adds less than n edges to it; thus the total number of added edges is less than mn yielding a complexity of $O(mn^3)$. □

It should be mentioned that we formulated Algorithm 2 as it was originally formulated by Lipton et al. However, one may interchange lines 5 and 6 and get the exact same result. The original version removes cycles when there is a need to while the modified version removes them the moment they are created. As a result, the modified version also maintains the invariant that E_G is a directed and *acyclic* graph (DAG); an observation that will be useful later on. Another useful observation is that in striking contrast to Greedy Round Robin, Envy Cycle Elimination removes the envy by removing the last added item, not the first.

4.2 State of the art results

The first positive results of the area are due to Plaut and Rougharden, [PR18]. Regarding exact EFX allocations, they showed that it is possible to compute one when the agents have general but *identical* valuations by introducing the *leximin++* operator. Leximin, a portmanteau of the words "lexixocgraphic" and "maximin", refers to the technique of picking some arbitrary agent ordering, e.g. agent 1 is more important than agent 2 etc, and maximizing the minimum value (that of agent n's bundle), then the second minimum and so on. Plaut and Roughgarden's tweak was that once the value of a bundle was maximal they also maximized its size. A small sketch proof: consider a non EFX allocation $\mathcal{A}$. Then there exists at least one pair of agents i, j that contradicts definition 3.3.2: $v(\mathcal{A}_i) < v(\mathcal{A}_j \setminus g)$ for some $g \in \mathcal{A}_j$. Construct a new allocation $\mathcal{B}$ by moving g to $\mathcal{A}_i$. Now, every agent less important than i has the same bundle and agent i has either $\mathcal{A}_i \cup g$ or some $\mathcal{A}_k$, which originally belonged to a more important agent. In either case, $\mathcal{B}$ is better than $\mathcal{A}$ with respect to the leximin++ operator. Therefore the leximin++ solution must be EFX.

Based on their solution for identical valuations they show that EFX allocations exist for two agents even with different valuations: one agent *cuts* the set of goods into two bundles based on her valuation and the other *chooses* her favorite. The *cut and choose* protocol, as it is named, is EFX because the cutter constructed the bundle based on the leximin++ solution therefore she is satisfied with both and the chooser picks the best available thus she is never envious.

In the restrictive domain of additive valuations, Plaut and Roughgarden showed that the identical valuations setting can be extended to identical rankings by running algorithm 2 with the goods in decreasing order of value. They were also the first to show how to compute an approximate EFX allocation, developing an algorithm for 1/2-EFX even for *subadditive*[1] valuations. Unfortunately, with the exception of identical additive rankings, every algorithm they presented is inefficient. They proved that this last caveat cannot be avoided when working past additive valuations.

Given the challenging nature of the problem, after the work of Plaut and Roughgarden, there was a shift of interest in improving the 1/2 approximation ratio. Chan et al., [CCLW19] presented a polynomial time algorithm that matches the approximation ratio. Their algorithm proceeds in rounds, computing a matching between agents and unallocated items in each round. In the next chapter we will adapt their proof technique to get a simpler algorithm by slightly tweaking the Envy Cycle Elimination algorithm. The first, and currently only, (unconditional) improvement upon the 1/2 approximation of [PR18] is due to Amanatidis et al., [AMN20]. Their algorithm outputs a $(\phi - 1)$-EFX allocation (where $\phi \approx 1.618$ is the golden ratio) and runs in polynomial time as well. The base of their algorithm is once again algorithm 2 equipped with a clever preprocessing step: the agents are partitioned in two sets, one where each agent receives her favorite and high valued good and the other where the agents receive two goods as a way of compensation.

Amanatidis et al. also introduced the concept of EFX *with few items*. When the number of items is not larger than the number of agents it easy to check that any allocation where no agent receives multiple goods is EFX. It is also trivial to find an EFX allocation when there are exactly $n + 1$ items, even if the agents have general valuations: the agents pick according to a given ordering, their favorite unallocated item with the last agent picking both the remaining goods. In [AMN20], it is shown how to get an EFX allocation with $n + 2$ items and additive valuations. We will show an elegant

[1] For two disjoint sets S, T a valuation is subadditive if $v(S \cup T) \leq v(S) + v(T)$

proof for extending this result to general valuations; however, Mahara already extended the result to $n + 3$ items in [Mah21]. Still, their analysis is cumbersome and spans multiple pages so our result may be useful in simplifying and/or further extending it.

With regard to exact EFX allocations, the first major breakthrough after the work of Plaut and Roughgarden is due to Chaudhury et al., [CGM20]: they demonstrated how to compute an EFX allocation for 3 agents with additive valuations. They studied the problem from a graph theoretical perspective, similar to Lipton et al., [LMMS04]; using the envy graph and the novel *champion graph*. Informally, the champion graph reflects, via its labelled edges, the largest envy after allocating some unallocated good. The result was later extended by Berger et al., [BCFF21], to the broader class of *nice cancelable* valuations. They showed that this new class of valuation functions exhibits some welcomed properties when chasing fair allocations which, in turn, lead us to studying them further in this thesis.

Definition 4.2.1 (Definition 2.1 in [BCFF21])**.** A valuation function v is cancelable if for any bundles $S, T \subset M$, and item $g \in M \setminus (S \cup T)$, it holds that

$$v(S \cup g) > v(T \cup g) \implies v(S) > v(T)$$

At the time of writing this thesis, a new result by Akrami et al., [ACG+22], was published which greatly extends the class of valuation functions that admit EFX allocations in the 3 agent setting. Namely, it suffices that only one agent's valuation abides by some mild condition while the other two agents can have general and possibly distinct valuations.

Beyond the study of approximate EFX allocations, and the settings of few agents or few items, there has been some progress in some more cases, although quite restricted as well. Aleksandrov and Walsh, [AW19], showed how to compute an EFX allocation for the class of binary additive valuations, Babaioff et al., [BEF21], for dichotomous submodular valuations and Amanatidis et al., [ABFR+21], designed algorithms for instances with 2 values ($v_i(g) = a$ or $v_i(g) = b$ for all goods and agents) and when all possible item values lie in an interval of the form $[x, 2x]$. The former result was extended by Gard and Murhekar, [GM21], where EFX was achieved in conjunction with Pareto optimality. The same conjunction was achieved by Hosseini et al., [HSVX21] for lexicographic preferences. As for the second result of Amanatidis et al., it will be later exploited in the design of algorithms with better approximation ratio. Finally, Mahara, [Mah20], initiated the study of allocations where every agent has one out of two possible valuations, proving that EFX allocations exist when the two valuations are additive. Later on the result was extended to nice cancelable valuations before it got completely settled, in [BCFF21] and [Mah21] respectively.

4.3 A new direction: EFX with charity

We will close this chapter studying a different relaxation of EFX. *EFX with charity*, as named by Caragiannis et al., [CGH19], relaxes the requirement of the allocation being complete; some items are left unallocated (donated to charity). Caragiannis et al. demonstrated that after starting with the optimal *Nash social welfare*[2] allocation and carefully discarding some items one can end up with an EFX allocation that enjoys at least half as much social welfare as the optimal. While the Nash social welfare is an important fairness measure in its own when dealing with multiagent problems, no more guarantees regarding exclusively EFX allocations were provided in [CGH19]. However, the idea of charity gained attraction after the work of Chaudhury et al., [CKMS21]. They introduced the following term.

Definition 4.3.1 (Bounded charity)**.** Let $\mathcal{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ and P be a partition of M into $n+1$ sets where $\mathcal{A}_i$ is the bundle of agent i and P is the bundle donated to *charity*. Moreover, we say that the charity is *bounded* if:

- $v_i(\mathcal{A}_i) \geq v_i(P)$ for every agent i,
- $|P| < n$.

They developed an algorithm that computes such an allocation with charity in pseudo-polynomial time (since there is a dependency on the value $v_i(M)$). They, also, showed how to modify it to obtain an FPTAS[3]. Here, we will just sketch the basic parts of the main algorithm.

The algorithm is, in a way, a sophisticated Envy Cycle Elimination. The change is that an item g is not allocated in every round as in line 4 of algorithm 2, but only if it does not disrupt the EFX property. This is update rule U_0, to stick with the original notation. The second update rule, U_1, deals with the possible envy towards the unallocated items, P. Chaudhury et al. introduced the notion of the *most envious agent*. Simply put, agent i is the most envious agent of a bundle X if the subset of X she envies has the smallest cardinality. For instance, if agent 1 envies $X \setminus g_1$ and agent 2 envies $X \setminus (g_2 \cup g_3)$, then agent 2 is more envious of X than agent 1. Update rule U_1 simply gives the most envious agent of P that minimal cardinality subset and adds her previous bundle to P. Update rule U_2 is the most complex one. We will explain it in two steps. Firstly, remember that the modified version of algorithm 2 maintains the envy graph as a DAG. Now, consider the case where there is exactly one source s in that DAG. Since U_0 is not applicable we deduce that some agent t is strongly envious of $\mathcal{A}_s \cup g$ and let Z be its minimum envied by s subset. The key observation is that t is reachable from s, i.e., there is a path of agents $s \to i_1 \to \cdots \to i_k \to t$. Changing s's bundle from A_s to Z

[2] The geometric mean of agents' valuations for their bundles

[3] *Fully polynomial time approximation scheme* is a solution concept in approximation algorithms. For more we refer the reader to [Vaz01]

adds the edge $t \to s$ to complete a cycle. Thus, we can move the bundles around ($\mathcal{A}_s = \mathcal{A}_{i_1}, \mathcal{A}_{i_{k-1}} = \mathcal{A}_{i_k}, \mathcal{A}_t = Z$) to decycle the graph as usual. Note that since Z was the envied subset of minimal cardinality no agent strongly envies it; otherwise $Z \setminus h$ would have been an envied subset of $X_s \cup g$ with even smaller cardinality. Therefore, the EFX property is maintained. Now consider that there are two sources s_1 and s_2 and t_1 and t_2 are the respective most envious agents of $\mathcal{A}_{s_i} \cup g_i$, $i = 1, 2$. If t_1 (resp t_2) is reachable from s_1 (resp. s_2) we can use the same approach as before. Assuming the contrary means that t_1 (resp. t_2) is reachable from s_2 (resp. s_1). Therefore, we can create the "merged" cycle $s_1 \to \cdots \to t_2 \to s_2 \to \cdots \to t_1 \to s_1$ and decycle it. Chaudhury et al. showed that as long as U_0 is not applicable and there are at least as many items in P as agent in $\mathcal{N}$ we can carefully craft a "merged" cycle. Finally, note that all three rules produce Pareto dominating allocations thus the algorithm is guaranteed to terminate.

The approach of bounded charity opened a new way to attack the problem: instead of a searching for a complete EFX allocation directly or gradually improving the approximation ratio, one can attempt to reduce the charity until no items get donated. It should be mentioned that both proofs for the 3 agent settings that were discussed in the previous paragraph, [CGM20] and [BCFF21] based their case analysis on allocating the 2 remaining items. Moreover, in [BCFF21] it was shown that for nice cancelable valuations one can reduce the size of charity to $n - 2$ item and, as with the case of two valuations functions, Mahara, [Mah21] extended the result to general valuations. The technique of charity also helped with one more result of [BCFF21]: in the case of 4 agents (with nice cancelable valuations) we can do one item better, leaving only one item unallocated.

The next step in reducing the number of donated items was, at least from a complexity theory standpoint, searching for sublinear charity. It was firstly achieved by Chaudhury et al., [CGM+21]. Specifically, they showed that an $(1-\epsilon)$-EFX allocation with high Nash welfare can be computed (in polynomial time) with charity $|P| \leq 64 \left(\frac{n}{\epsilon}\right)^{4/5}$. Interestingly enough, they obtained the result via a connection with extremal graph theory and combinatorics. Said connection allowed the independent study of the problem and led to further improvement by Berendsohn et al., [BBK22], where the number of items was reduced to $O(n^{2/3})$. The same complexity result was also obtained in [ACG+22]. Since the connection with the field of combinatorics was established, both works improved upon bounds from the works of Alon and Krivelevich, [AK21], and Meszaros and Steiner,[MS21], on zero-sum combinatorics.

The technique of bounded charity has also started being applied in restricted settings. Akrami et al., [ARS22], presented an algorithm that produces an exact EFX allocation, i.e., not $1-\epsilon$, with less than $n/2$ unallocated items when the agents have *restricted additive* valuations, a class that extends identical additive valuations allowing some agents to not value some good(s).

Chapter 5

Contribution

In this chapter we present the theoretical contributions of the thesis. The first part of the chapter is dedicated to additive valuations where we show how to obtain better approximations under various conditions. In the second part, we generalize the result of Plaut and Roughgarden, [PR18], from identical additive valuations to identical cancelable ones and introduce the notion of a ranking with tiers as a further extension. Our work is based on case analysis on the number of sources in the envy graph. Finally, we show that the same technique can be applied even in some settings with general valuations to simplify known results.

5.1 Approximations for additive valuations

In this paragraph, we attempt to go beyond the $\phi-1$ approximation ratio of Amanatidis et al., [AMN20]. While we did not manage to obtain a general result for all additive valuations, we will show some approaches that work under certain conditions.

5.1.1 Preferential Envy Cycle Elimination

We begin with a new way to obtain a 1/2-EFX in polynomial time. While the result is not interesting in itself, since there is a known 1/2-approximation, we will then build upon our proposed algorithm. The algorithm is quite simple, and consists of running one round of Greedy Round Robin (GRR from now on), followed by Envy Cycle Elimination (ECE).

Algorithm 3 Envy Cycle Elimination with a top preference

1: Run one round of algorithm 1 (GRR)
2: Continue with algorithm 2 (ECE) until there are no unallocated items

Note that algorithm 3 is equivalent with running ECE with the tweak that in the first n iterations (where there always exists at least one source), we first select a source agent and then have the agent select her favorite item, hence its name. We are now ready to prove our first result.

Theorem 1. Algorithm 6 computes an $\frac{1}{2}$ EFX allocation in polynomial time.

Proof. First, note that if the number of items is at most n, the algorithm is trivially EFX. Hence assume $m > n$ for the sequel. Let M_1 be the set of unallocated goods after the execution of the single round robin round. By that time, each agent i has received exactly one good g_i and the allocation is trivially EFX. Since $m > n$, we know that $M_1 \neq \emptyset$, and we have that

$$\forall i \; \forall h \in M_1 : \; v_i(g_i) \geq v_i(h)$$

Since the second step of the algorithm can never decrease an agent's valuation, then if $\mathcal{A}_i$ is the bundle of agent i during any phase of the algorithm, it holds that

$$\forall i \; v_i(\mathcal{A}_i) \geq v_i(g_i) \implies \forall h \in M_1 : \; v_i(\mathcal{A}_i) \geq v_i(h) \tag{5.1}$$

Since the initial allocation of the round robin step is EFX, and thus EF1, the final allocation will also be EF1 as step 2 maintains the EF1 property. Fix agent i and assume that she envies another agent j. Then, by the EF1 property it holds that

$$v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j \setminus g') \tag{5.2}$$

where g' is the last item added to $\mathcal{A}_j$ (which may not have belonged to agent j at the time). We can assume that $g' \in M_1$ (otherwise, $|A_j| = 1$ and agent i trivially satisfies the EFX property w.r.t. j). Using (5.1) for $h = g'$, and adding it to (5.2) gives

$$2 \cdot v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j \setminus g') + v_i(g') = v_i(\mathcal{A}_j)$$

The last relation implies that the allocation is at least 1/2 envy free (1/2-EF) and thus 1/2-EFX as claimed. Given than both procedures run independently and in polynomial time the proof is completed. □

Remark 1. We used the term 1/2 envy free without having formally defined the approximation version of envy freeness. Still, it works like Definition 3.4.1: $v_i(\mathcal{A}_i) \geq \alpha \cdot v_i(\mathcal{A}_j)$. Most results in this section imply an α-EF allocation when there are enough items but we will refrain from mentioning it.

The same result is obtained via a different algorithm in [CCLW19]. We should also note that the result of Theorem 1, in contrast to the next ones, holds even for subadditive valuations. So we just saw how adding preferential selection only in the first step of ECE improves the resulting allocation from EF1 to EF1 and 1/2-EFX. The next logical step is to check what happens when we do the same modification in every step of the way. The resulting algorithm, as presented below, was actually introduced by [BBKN18] with a different aim in mind.

Algorithm 4 Preferential Envy Cycle Elimination

```
1: Set 𝒜_i = ∅ for every agent i
2: Let M' = M denote the currently unallocated items
3: while M' ≠ ∅ do
4:     if ∃ some source s then
5:         Set 𝒜_s = 𝒜_s ∪ arg max_{g'∈M'} v_s(g')
6:         Update M'
7:     end if
8:     Decycle the envy graph
9: end while
10: return 𝒜 = (𝒜_1, ..., 𝒜_n)
```

Line 5 is the single change between algorithms 2 and 4. Furthermore, this is the algorithm Plaut and Roughgarden proposed for the setting with identical rankings in [PR18].

Theorem 2. Let E_G denote the resulting envy graph after running Preferential Envy Cycle Elimination (pECE). For every directed edge $(i, j) \in E_G$ let t_{ij} be the number of times agent i got to pick a good before the edge was created and set $k = \min_{i,j} t_{ij}$. Then, pECE computes a $\frac{k}{k+1}$-EFX allocation.

Proof. Let $\mathcal{A}$ be the allocation output by the algorithm, and fix some agent i that is envious of another agent j in $\mathcal{A}$. Let $k_i = \min_i t_{ij}$ for some edge e. Note that agent i's bundle at the time e gets added may not be the collection of the items she picked, C_i, due to some envy cycle elimination(s). However, since the algorithms proceeds without decreasing the valuation of any agent, we can guarantee that i prefers her bundle at the time, say T_i, to C_i. Since agent i was always picking her favorite available item, it holds that

$$\forall g \in C_i \ \forall h \in M' : \ v_i(g) \geq v_i(h)$$

By combining the inequalities above, we deduce from the additivity of v_i that

$$v_i(T_i) \geq v_i(C_i) \geq k_i \cdot v_i(m) \tag{5.1'}$$

Equation (5.1′) substitutes (5.1) of the previous proof. Similarly we have

$$k_i \cdot v_i(T_i) \geq k_i \cdot v_i(T_j \setminus g') \tag{5.2'}$$

where we simply multiplied by k_i. Once again setting $h = g'$ and adding yields

$$v_i(T_i) \geq \frac{k_i}{k_i + 1} v_i(T_j)$$

Note that once someone envies j, she will never become a source, therefore $T_j = \mathcal{A}_j$, and $v_i(\mathcal{A}_i) \geq v_i(T_i)$ meaning that $v_i(\mathcal{A}_i) \geq \frac{k_i}{k_i+1} v_i(\mathcal{A}_j)$. Thus the allocation is $\frac{k}{k+1}$-EFX where $k = \min k_i$ and the proof is completed. □

Clearly, everybody gets to pick at least once, thus $k \geq 1$. Hence, this is at least as good as the $1/2$ bound established earlier, but is it always strictly better? The following example shows than in worst case, this algorithm still gives only a $1/2$-approximation.

Example 5.1. Consider the following instance with 2 agents and 4 items where every tie is broken lexicographically:

	a	b	c	d
agent 1	2	2	0	2
agent 2	2	1	$1-\epsilon_1$	ϵ_2

with $0 < \epsilon_1 < \epsilon_2 < 1/2$. The algorithm gives a to agent 1 and the rest items to agent 2. It is easy to verify that the allocation is exactly $1/2$-EFX. Moreover, the issue is not the number of agents since we can add as many single minded[1] agents and one item for each. △

However, not all hope is lost. The analysis of algorithm 4 tells us that the more items an agent receives, the better the approximation which, as a first step, leads to the following corollary.

Corollary 2.1. Assuming some sort of large market where each agent must receive at least l goods, e.g., every bundle of size l is valued at most at $\epsilon\dfrac{v(M)}{n}$ for some $\epsilon \ll 1$, then there exists an $\frac{l}{l+1}$-EFX approximation.

5.1.2 Relaxed condition-approximation tradeoffs

The previous corollary may seem a bit vague, but it points us to the right direction. The troublesome cases in the pursuit of EFX, at least when agents have additive valuations, are when the agents value the top items much more than the rest. In order to make progress we may try to condition our instances accordingly. One such simple scenario to consider is when agents have different top preferences.

Theorem 3. Assuming that each agent i has a different favorite good h_i then a $2/3$-EFX allocation can be computed efficiently.

Firstly, we show how to properly modify algorithm 3 and then we prove the claim.

Algorithm 5 Envy Cycle Elimination with different favorite items

1: Run two rounds of algorithm 1 (GRR)
2: Continue with algorithm 2 (ECE) until there are no unallocated items

Proof. Let $S_i = \{h_i, g_i\}$ be the bundle of agent i after the execution of step 1. Then, for any pair of agents i and j,

$$v_i(S_i) \geq v_i(h_i) > max_{g\in\{h_j,g_j\}} v_i(S_j \setminus g)$$

[1]Agents interested in a single good

Or, in words, agent i satisfies the EFX condition w.r.t. the other bundles of size 2. Thus the initial allocation S is EFX. Moreover,

$$\forall h \in M \setminus \bigcup_{i \in [n]} \{h_i, g_i\} : \ v_i(S_i) = v_i(h_i) + v_i(g_i) > 2v_i(h) \tag{$*$}$$

And the analysis is now identical to that of pECE with $k = 2$. □

At that point, one may argue that the condition is too weak to have any practical value. We will return to it later but we will argue that it gives a hint on the needed conditions to improve the approximation ratio. The intuition behind Theorem 2 was that you need a lot of goods to guarantee good approximations. Theorem 3 shows us that it suffices to force an initial allocation with few items but not singleton sets. In reality, what we really need is some property like in Equation ($*$). Going back to the proof of Theorem 2 that was equation (5.1′). Putting everything together we form the following general framework:

Algorithm 6 General approximation framework

1: Start with the empty allocation
2: Compute a partial α-EFX allocation S maintaining the property

$$v_i(S_i) \geq \beta \cdot v_i(h) \text{ for all } h \in M \setminus S$$

3: Continue with algorithm 2 (ECE) until there are no unallocated items

Theorem 4 (Approximation Framework)**.** Algorithm 6 computes a $\min\left(\alpha, \frac{\beta}{\beta+1}\right)$-EFX allocation. Moreover, whenever the partial allocation can be computed efficiently, the same holds for the whole allocation of the algorithm as well.

Under our approximation framework, the algorithm of Amanatidis et al., [AMN20], achieves its $\phi - 1$ ratio by setting $\alpha = \beta = \phi$.

Up to this point we have not discussed anything about relaxed conditions. Notice that our general approximation framework is plug and play: if you have an algorithm for step 2 you are good to go. Now, recall that the first result about additive valuations is under the common item ranking assumption due to Plaut and Roughgarden, [PR18]. What if we relaxed that condition and only ask for a common ranking of the most important items?

Corollary 4.1 (Relaxed top ranking)**.** Assuming that all agents agree upon the order of the top l items, then one can efficiently compute a $\frac{k}{k+1}$-EFX allocation with $k = \lfloor l/n \rfloor$.

Proof. For $l < 2n$ we do not improve upon the bound of $1/2$ anyway, thus we will prove the claim for $l \geq 2n$. By [PR18], running the pECE algorithm for the top l items constructs a partial EFX allocation ($\alpha = 1$). For an agent i with a bundle of size at least k it obviously holds that $v_i(S) \geq k \cdot v_i(h)$ ($\beta = k$). If there is some agent j with a bundle of size less than

k then, by the pidgeonhole principle, there exists some j' with at least $k+1$ items. The EFX property of S gives $v_j(S_j) \geq v_j(S_{j'} \setminus g)$, where the last bundle is of size at least k, and therefore $v_j(S_j) \geq kv_j(h)$. Thus we have everything the general framework requires. □

We provide yet another corollary of the approximation framework.

Corollary 4.2 (Relaxed top bounded interval)**.** Assuming that all agents value the top l items between x and $2x$, then one can efficiently compute a $\frac{k}{k+1}$-EFX allocation, with $k = \lfloor l/n \rfloor$

Proof. The proof is the same as above except that S is produced by algorithm 2 of [ABFR+21]. □

Similarly, one can obtain an approximate EFX allocation by relaxing any constrained domain where a full EFX allocation is known to exist to the first few items. Still, the main question on the agenda of EFX approximations is achieving the ratio of 2/3.

5.1.3 2/3 approximation for common top n rankings

The result of the previous paragraph means that a 2/3-EFX allocation exists when the agents have a common ranking of the top $2n$ items or when their top $2n$ items are valued within the interval $[x, 2x]$. It turns out we can do even better. Before presenting the algorithm, let us build some intuition first. Ideally, we would like to run two complete rounds of GRR, like in algorithm 5. That is too much to ask for. A more reasonable alternative would be to run one round of GRR and the second in reverse order. Clearly, this procedure is fairer in the sense that agents with a bad pick in round one may offset it in round two. Unfortunately, it is far from certain than an agent with two goods will not envy one with a single item.

In algorithm 7 we treat this problem by simply giving to the agent with two items the single one she envies the most, which is practically equivalent to restarting the process after putting said agent higher in the order. Note, however, that in contrast with the previous algorithms, algorithm 7 may (temporarily) unallocate some goods (the old bundle of i before line 7).

Theorem 5. Assuming that all agents agree upon the ranking of the top n items, then algorithm 7 computes efficiently a 2/3-EFX allocation.

Proof. Firstly, note that if $m \leq n$, the allocation is EFX by step 2 alone and no further work is needed. For the rest of the proof we will assume that $m > n$. Secondly, we will show that when the loop terminates (or when we run out of items in its middle) the partial allocation, say S, is EFX. To do so we will start by showing that every time the loop starts all of the top n items are allocated. That holds trivially the first time. If $i = j$ no unallocations happen so it still holds and when $i \neq j$ agent i returns her item, say g_i, to get j's top n item (line 7). Then agent j picks g_i (line 8) since it is the only top n item

Algorithm 7 2/3 EFX for identical top n rankings

1: Pick an arbitrary ordering of the agents π: $\pi(i) = j$ means that i is in the j-th position
2: Run one round of algorithm 1 (GRR)
3: Let M' be the unallocated items after the GRR round
4: **for** i in range(N,1) **do**
5: Set $g' = \arg\max_{g \in M'} \; v_i(g)$
6: Set $\mathcal{A}_i = \mathcal{A}_i \cup g'$
7: Select agent $j = \arg\max\limits_{k \in [i]} \; v_i(\mathcal{A}_k)$
8: **if** $i \neq j$ **then**
9: $M' = M' \cup A_i$
10: $\mathcal{A}_i = \mathcal{A}_j$
11: $\mathcal{A}_j = \arg\max \; v_j(m)$
12: Shift every agent from $i-1$ to $j+1$ one position forward
13: Set j in the front $\pi(i-1) = j$
14: **else**
15: $M' = M' \setminus g'$
16: **end if**
17: **end for**
18: Continue with algorithm 2 (ECE)

available. In other words, i and j simply swapped their initial items. In any case, any agent has exactly one top n item and maybe one extra.

To see that the EFX property is maintained pick agent i in the ℓth position of the final ordering. Any agent j in position $\ell' > \ell$ either has exactly one item thus EFX is trivial or has two items but she picked her top n item $g_{\ell'}$ after the agent in position ℓ. Thus $v_i(S_i) \geq v_i(g_{\ell'}) = v_i(S_j \setminus g)$. On the other hand, the quantifier in line 5 informs us that agent i does not envy any agent with $\ell' < \ell$, when she is in the front, therefore she cannot strongly envy them after they receive one more good.

To complete the proof we need the relation

$$v_i(S_i) \geq \beta \cdot v_i(h) \text{ for all } h \in M \setminus S$$

for $\beta = 2$. If i has a bundle of size 2 then the property holds trivially after her turn is completed. It suffices than no items get unallocated after. Since the main loop does not permanently unallocate the top n items we are done. If agent i ended up with only one item, say a, then we can write

$$v_i(a) > v_i(b) + v_i(c) \geq 2 \cdot v_i(c)$$

where $\{b, c\}$ is her bundle before the check $i \neq j$ with b the top n item. So it suffices for b and every item in front to not get unallocated but again those are top n items. Thus algorithm 7 matches the approximation framework. Since every step of the main loop is at most linear in n and m the algorithm is also efficient. □

The logical next step is to check what the algorithm can do without the common top n ranking assumption.

Corollary 5.1. Allowing a little charity (up to $n-1$ items), a 2/3-EFX allocation can be computed efficiently.

Proof. To achieve such an allocation we will use a slight modification of algorithm 7; during the main loop every agent will mark the items she values more than half of her current bundle. Then, before running Envy-Cycle-Elimination all the marked items will be donated to charity. Firstly, note that the 2/3 approximation follows trivially from the previous proof since all the items disrupting the property then were top n while now are donated. It remains to show the bound on the number of items. To that end, note that agent n will mark at most $n-1$ items. After that, every time a marked good (g_i) is unallocated (line 7) and replaced by a new one to be potentially marked (line 8), one marked item (g_j) is guaranteed to remain in the final allocation, thus the number of marked goods can never exceed $n-1$. □

A final remark before leaving additive valuations behind us: when $\alpha = 1$ the approximation framework actually guarantees something stronger than $\frac{\beta}{\beta+1}$-EFX. It gives an allocation that is either EFX or $\frac{\beta}{\beta+1}$-EF. As a result, we believe that one can further weaken the top n ranking condition, or eliminate it altogether to achieve the desired 2/3 approximation.

5.2 Tiers for cancelable valuations

We now shift our focus to cancelable valuations. As mentioned in Chapter 4, a version of them was introduced by Berger et al., [BCFF21] where it was shown that many results that hold for additive valuations can be extended to this broader class. Before adding one more result to this list, we will offer an explanation as to why cancelable functions seem to behave well. Recall that in Definition 4.2.1 the bundle $S \setminus g$ is mentioned. That is the same as in the definition of EFX (3.3.2). Perhaps, the class of cancelable valuations is the broader one that encompasses such information into its definition.

5.2.1 Common ranking

Without any further delay, we are ready to state and prove the first result of this section, which is a generalization of [PR18] under the common ranking assumption on the goods:

Theorem 6. Algorithm 4 (pECE) efficiently computes an EFX allocation when all agents have a common ranking of all goods and cancelable valuations.

Proof. Let $g_1, \ldots, g_m$, denote the ordering of the goods. Obviously, this is also the order in which the goods get allocated by the algorithm. Assume that after the allocation of

g_k some agent i strongly envies agent j who received it. Here, we remind the reader the discussion in Chapter 3: strong envy means that the EFX property is violated. Then it holds that

$$v_i(\mathcal{A}_i) < v_i(\mathcal{A}_j \cup g_k \setminus g_l) \text{ for some } l < k$$

Since j received the good, she was previously a source, that is, unenvied:

$$v_i(\mathcal{A}_i) \geq v_i(\mathcal{A}_j)$$

Combining the two inequalities gives

$$v_i(\mathcal{A}_j \cup g_k \setminus g_l) > v_i(\mathcal{A}_j)$$

And by iteratively applying the definition of cancelability for any good in $\mathcal{A}_j \setminus g_l$ yields

$$v_i(g_k) > v_i(g_l) \text{ for some } l < k$$

which contradicts the common ordering. □

We must note that the same result was obtained independently by Garg and Sharma, [GS22]. However, after this point the two works follow a completely different path.

5.2.2 Tiers of size 3

We will now proceed to a relaxation of the common ranking assumption. Our path starts with a technical lemma.

Lemma 1 (Inequalities under addition)**.** Let S, T, Q and R be sets such that $S \cap Q = \emptyset$, and $T \cap R = \emptyset$. Then

$$\left.\begin{aligned} v(S) \geq v(T) \\ v(Q) \geq v(R) \end{aligned}\right\} \implies v(S \cup Q) \geq v(T \cup R)$$

Proof. In the definition of cancelable valuations, it is easy to see that one direction implies the opposite: $v(T) \leq v(S) \implies v(T \cup g) \leq v(S \cup g)$. Applying this form of the definition for every $g \in Q \setminus T$ gives

$$v(T \cup Q) = v(T \cup (Q \setminus T)) \leq v(S \cup (Q \setminus T)) \leq v(S \cup Q)$$

where the last inequality is due to the monotonicity assumption. Similarly we obtain $v(Q \cup T) \geq v(R \cup T)$ and the lemma follows.

□

Before delving into the details about the new notion of *tiers* that we are about to introduce, let us consider a quick warm-up. We just saw that when the rankings are common, an EFX allocation exists for cancelable valuations. Previously, we saw also that only a common ranking of the top n items suffices for a good approximation. Can we somehow combine the two settings? For example, what if the agents agree on the set of the best n items (but not necessarily on their ordering), and then agree on the set of the second best n items, and so on? Those settings are common in real life: it is often the case that when asked to do a list of favorite athletes, movies, etc., people find it too hard to produce a complete ordering and instead use *tiers* to denote their preferences. In such situations, several people may agree on the elements contained in each tier, but without necessarily agreeing on the ordering within the tiers.

Definition 5.2.1. A partial ranking $\mathcal{T}$ is a common *tiered* ranking among all agents, if there exists an ordered partition of all items $M = (M_1, M_2, \ldots, M_l)$, such that:

$$\forall g \in M_k, \forall h \in M_{j>k} : \; v_i(g) \geq v_i(h)$$

for every agent i. Moreover, we define the size of $\mathcal{T}$ to be the size of the largest tier M_j.

Based on the above definition the question raised earlier can be restated as: does an EFX allocation exist when all agents have common tiered rankings of size n? Unfortunately, we failed to answer this. Still, some measurable progress was made. Note that the common ranking is equivalent to a tier of size 1. A natural way to progress is to try increasing the size to 2, 3 etc. The main result of this paragraph is the following:

Theorem 7. Assuming that all $n \geq 3$ agents have cancelable valuations and a common tiered ranking of size at most 3, i.e.

$$a_1, b_1, c_1 \succeq a_2, b_2, c_2 \succeq \ldots \succeq a_k, b_k, c_k$$

then an EFX allocation exists and can be computed efficiently.

Proof. We will prove the theorem by induction. We present the proof for the case where each tier has exactly 3 items, since the other cases are easier. Note that the base of our induction, allocating the first triplet, is easy: just give a single item to three different agents. For the inductive step assume that we have a partial EFX allocation after allocating all the goods up to some tier k, and let E_G be the envy graph of this allocation. Moreover, assume that we maintain the envy graph E_G of this current allocation as a DAG. If it is not, we can always remove all the envy cycles prior to continuing with the next tier. Also, we will simplify the notation from a_{k+1} to just a and so on. We discern three cases based on the number of sources in E_G, and discuss them in order of difficulty.

- Case 1: E_G has at least three sources

  This is the easiest case since we have three items and we can pick three sources and just do a matching. The EFX property is maintained and the proof is analogous to that of Theorem 6.

For the two remaining cases we will have less sources than items. Let a, b, c be the 3 items of the tier under consideration, and let $(\mathcal{A}_1, \ldots, \mathcal{A}_n)$ be the current EFX allocation of the goods in the first k tiers. On a high level, both case analyses work as follows: the agents closest (in the sense of topological distance) to the sources are possible new sources. Therefore, if our current sources outvalue them or maybe get their bundles via some envy cycle elimination we can allocate the remaining items. Otherwise, some current source will receive more than one items without violating the EFX property. Some extra notation: we will refer to agents based on their level when we view E_G as a DAG that is topologically ordered, i.e. all sources are at level 0, agents envied by the sources at level 1 etc.

- Case 2: E_G has one source s_1

  Let $o_1, o_2, \dots, o_\ell$ be the agents of level 1 ordered based on s_1's valuation: $v_{s_1}(\mathcal{A}_{o_1}) \leq v_{s_1}(\mathcal{A}_{o_2}) \leq \dots \leq v_{s_1}(\mathcal{A}o_\ell)$.

  - Subcase 2a: s_1 can receive multiple items

    If s_1 can receive all three goods of the tier without violating the EFX property we are done. If not but she still can receive two goods, say a and b, we allocate them to her. It remains to allocate c. If $v_{s_1}(\mathcal{A}_{s_1} \cup a \cup b) \geq v_{s_1}(\mathcal{A}_{o_1})$ we allocate c to o_1 and complete the inductive step. Indeed, note that c is the least valuable item of the bundle that o_1 has now. Since s_1 was not envying o_1 after she received a and b, the EFX property will not be violated by giving c to o_1. There is no other pair of agents that we need to check since no one else is allocated any items. Suppose now that $v_{s_1}(\mathcal{A}_{s_1} \cup a \cup b) < v_{s_1}(\mathcal{A}_{o_1})$. We also know that since s_1 could not receive all 3 items, some agent, say x, must envy her. The fact that s_1 was the single source of E_G means that x is reachable from s_1 via an envy path. Therefore the allocation of a and b to s_1 has created an envy cycle $s_1 \to o_i \to \dots \to x \to s_1$, for some agent o_i of level 1. After decycling the graph, s_1 will be in possession of the bundle $\mathcal{A}_{o_i}$. Since previously she was the only one envying the bundle, we can now allocate c to her without disrupting EFX.

  - Subcase 2b: s_1 cannot receive multiple items If allocating 2 items is problematic for the EFX property, we deduce that after giving, say a, to s_1, some agent x becomes envious of her. Firstly, we will identify possible new sources. Those are the o_i agents, and more specifically o_1 and o_2, and agents envied only by o_1 (and maybe s_1), which we denote by $t_1, t_2, \dots$. The nodes of interest are shown in Figure 5.1.

    Similarly to the previous case, if $v_{s_1}(\mathcal{A}_{s_1} \cup a) \geq v_{s_1}(\mathcal{A}_{o_2})$, we allocate one item to both o_1 and o_2 and we are done. Otherwise, we continue in the same spirit as before. Suppose first that x is reachable from o_r, for some $r \geq 2$. Then, given that $v_{s_1}(\mathcal{A}_{s_1} \cup a) < v_{s_1}(\mathcal{A}_{o_2}) \leq v_{s_1}(\mathcal{A}_{o_r})$, an envy cycle $s_1 \to o_r \to \dots \to x \to s_1$ is created. After decycling it, we allocate one item to o_1 and one to the current owner of o_2 (it could be either o_2 or s_1), and we are done. And this is where

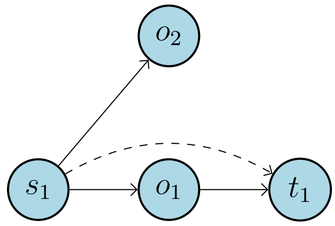

Figure 5.1: Subcase 2b

the similarities between the two subcases stop since x could be some agent only reachable from o_1 or she could be o_1 herself. We need to look at these two further cases separately but before doing so, notice that for agents who are unreachable from o_1, whoever owns $\mathcal{A}_{s_1}$ can get two items.

* Subcase 2b(i): $r = 1$, and x can be some agent other than o_1

  In this subcase, the role of the t_i nodes becomes more clear. So far, the possible new sources were always o_1 and o_2. Now they are o_1 and some node t_i. To proceed, note that there may not be a path from s_1 to x after allocating item a to s_1, but we will reallocate the bundles as if there was; checking that the EFX property is maintained is easy. Also, we pick x as the agent furthest away from s_1. The image looks as follows:

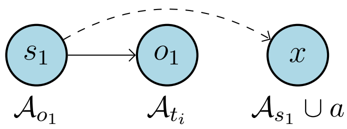

Figure 5.2: Subcase 2b(i)

  Now s_1 with her new bundle is again a source (single if x is not one) and if any node other than x is envious of $\mathcal{A}_{o_1} \cup b$ or $\mathcal{A}_{o_1} \cup c$ the next source will be either o_1 (owning $\mathcal{A}_{t_i}$) or $o_i, i \geq 2$ and we are done. The same applies if s_1 can stop envying some t_i after she receives a good since said t_i will become the final source or if s_i will receive $\mathcal{A}_{t_i}$ after some envy cycle elimination. If neither is true and s_1 cannot receive both of the remaining items without violating the EFX property, x must be a source as well and some agent reachable only from her envies s_1 after she receives a good. In this scenario, if one of the two matchings between the two sources and the two items produce an EFX allocation, we have completed this case. Otherwise, there is a cycle containing both sources and nodes from one or both connected components (Lemma 1 guarantees that no agent between s_1 and x can strongly envy the other one) if we substitute a with b or c.[2] Now, s_1 is in possession of some $\mathcal{A}_{t_i}$ and she can have the last item.

[2] Remember the discussion in paragraph 4.3. Now $Z = \mathcal{A}_{s_1} \cup b$ or $\mathcal{A}_{s_1} \cup c$

- Subcase 2b(ii): $r = 1$, and x is o_1

  In this last case, we have only one source and one possible new source. However, apart from o_1, no other agent would strongly envy $\mathcal{A}_{s_1} \cup a \cup b$ (or any other combination of $\mathcal{A}_{s_1}$ with a pair of goods from the given tier); otherwise we would be back to Subcase 2b(i). Therefore, we will compensate the lack of possible sources by possibly adding two items to one bundle. We start by asking s_1 to choose between receiving her favorite item or o_1's favorite, thus creating an envy cycle of size 2 and causing a swap. Note, however, that if s_1 and o_1 have a different favorite item then s_1 will always choose to swap since she well get her favorite item right after and lemma 1 guarantees the optimality of the choice. At any case, and with a the favorite item of o_1, the owner of $\mathcal{A}_{o_1}$ will be the new source. If she can get both of the remaining items the proof is completed. Otherwise some agent envies $\mathcal{A}_{o_1} \cup b$ or $\mathcal{A}_{o_1} \cup c$. Since we have a single source that envious agent is reachable so we ask our source to choose between her favorite item or the one that forms the envy cycle implied above. At any case, the new source will be eligible to receive the last item even if her bundle is $\mathcal{A}_{s_1} \cup a$.

- Case 3: Two sources s_1 and s_2

  Now, we have more sources but also a harder time identifying the possible new one. To bypass this problem we partition the envy graph E_G in the following manner:

  $$E_G = s_1 \cup s_2 \cup V_1 \cup V_2 \cup V_{12}$$

  where V_1 (resp. V_2) is the set of nodes reachable *only* from s_1 (resp. s_2) via an envy path and V_{12} is the set of nodes reachable from both sources. Since E_G is a DAG it follows that V_1, V_2 and V_{12} are as well. If s_1' is a source of the V_1 DAG we have that s_1 is the only agent envious of her; otherwise it would be reachable from s_2 thus contradicting the definition of the partition. Therefore, s_1' is a possible new source substituting s_1 and, symmetrically, a source s_2' of V_2 is a candidate substitution for s_2. A source s_{12}' of V_{12} may be a possible replacement for both. Now, if s_1 can receive two items or one and simultaneously stop envying s_1' we easily allocate all the items of the tier. Therefore, we assume this is not the case and denote with e_1 her envious agent after receiving some item of the current tier, and respectively, e_2 for s_2. We will do some case analysis based on which connected component e_1 and e_2 belong to. Since they may not be unique, we define T_1 to be the set of those envying s_1 (and respectively for s_2). Fortunately, due to symmetry the number of different subcases is small.

  - Subcase 3a: $T_1 \cap V_1 \neq \emptyset$ (or resp. $T_2 \cap V_2 \neq \emptyset$)

    Let $e_1 \in T_1 \cap V_1$. In that case e_1 is reachable from s_1 and we have an envy cycle. After applying a decycling step, s_1 will own $\mathcal{A}_{s_1'}$, a bundle of which she previously was the only envious agent. Therefore, we can allocate one more item to s_1 and the other to s_2 and we are done.

- Subcase 3b: $T_1 \cap (s_2 \cup V_2) \neq \emptyset$ (or resp. $T_2 \cap (s_1 \cup V_1) \neq \emptyset$)

  Let $e_1 \in s_2 \cap V_2$. Now, e_1 is not reachable from s_1. However, the same must apply to s_2 and e_2 otherwise the symmetrical of subcase 3a holds for s_2. In other words, $e_2 \in E_G \setminus V_2$. This means that, after allocating two of the three items of the tier, a cycle in the form: $s_1 \to \cdots \to e_2 \to s_2 \to \cdots \to e_1 \to s_1$ is created. Once we decycle the graph, checking that the EFX property is maintained is trivial. Moreover, there will be again two sources. If one of them does not own one of the two bundles previously owned by the sources, we can allocate the final item. In the unique case where $e_1 = s_2$ and $s_1 = e_2$ were the only choices for envious agents, we can allocate the last item to any of them. To see why, assume that s_1 owns $\mathcal{A}_{s_2} \cup a \cup c$. Agent s_2 cannot strongly envy her and if some other agent x does it means that $e_1 = s_2$ was not the only choice.

After careful inspection, subcases 3a and 3b and their symmetrical cover for 8 out of the 9 possible scenarios. It remains to check when both envious agents are in V_{12}.

- Subcase 3c: $T_1 \cap V_{12} \neq \emptyset$ and $T_2 \cap V_{12} \neq \emptyset$

  Let $e_1, e_2 \in V_{12}$. Note that since e_1 is reachable from s_1 if the first node in the path between them belongs to V_1 the argument of subcase 3a applies. Therefore, the first envy edge in the path is from s_1 to some source of V_{12}. If e_1 and e_2 belong to different weakly connected components of V_{12} we are done since one of the owners (s_1 or s_2) of some $\mathcal{A}_{s'_{12}}$ will get the last item. Assuming the contrary, the image is given in Figure 5.3.

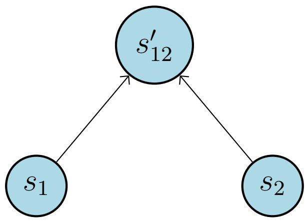

Figure 5.3: Subcase 3c

To continue, we select the source, say s_1, and the item, say a to allocate based on e_1's preference, who we pick to be in maximum topological distance. (Determining the largest distance can be achieved by allocating all items to s_1 one by one and checking) After reallocating the bundles along the cycle (if s_1 stops envying s'_{12} she becomes a possible new source similar to case 2 and we finish accordingly) and the image changes as in Figure 5.4.

The way we picked e_1 she and every possible node reachable from her cannot strongly envy $\mathcal{A}_{s_2} \cup b \cup c$; otherwise the strongly envious agent would have been

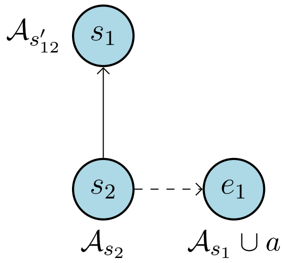

Figure 5.4: Subcase 3c continued

envious before and in a greater topological distance[3]. Therefore, either s_2 will get both items or there will be an envious agent she can reach. In the end, whoever owns $\mathcal{A}_{s'_{12}}$ will get the last item and the proof is completed.

□

5.3 Simplifications for general valuations

We have just seen how case analysis based on the number of sources in the envy graph can produce some positive results. With general valuations the analysis could be even harder so we will focus only on instances with a small number of them (one or two). Our work is once again built upon the ECE algorithm, equipped with a new allocation rule which combines those introduced by Chaudhury et al., [CKMS21]. If U is the set of currently uncallocated then our update rule allocate to some source s the minimum envied subset of $\mathcal{A}_s \cup U$. By construction of the rule, which we will call $\mathcal{U}$, no strong envy towards s can be created.

5.3.1 EFX with identical preferences

We start with an alternative proof about the existence of EFX allocations with identical but general valuations, [PR18]. Actually, we prove a slightly stronger claim: EFX allocations exist when agents have identical preferences; for every two subsets S, T of $\mathcal{M}$ all agents prefer the same set.

Theorem 8. When all agents have general valuations but identical preferences an EFX allocaton can be computed by running ECE with the $\mathcal{U}$ rule.

Proof. Let $\mathcal{A}$ be the allocation at some step of the algorithm. Since the agents have identical preferences $v_s(\mathcal{U}(\mathcal{A}_s)) > v_s(\mathcal{A}_s)$ thus the new allocation Pareto dominates $\mathcal{A}$ which guarantees the termination of the algorithm. □

[3] To be precise, e_1 could have been envious of $\mathcal{A}_{s_2} \cup b$ but then we restart the tier allocation working with s_2

5.3.2 EFX for two agents

In the previous section there was always one source. Moving on to the next simpler setting, we present a non *cut and choose* proof for the existence of EFX allocations for two agents. Our algorithm is again the same, with the difference that whenever the allocation is envy free with deterministically select agent 2 as the source.

Proof. Consider the following potential function

$$\phi(\mathcal{A}) = \big(v_1(\mathcal{A}_1), v_1(\mathcal{A}_2)\big)$$

- Agent 1 is the source
  Then either $\mathcal{A}_1$ will be directly improved from the application of $\mathcal{U}$ or mutual envy will occur resulting in a swap. Since $v_1(\mathcal{A}_2) > v_1(\mathcal{A}_2), \phi$ increases.

- Agent 2 is the source or the allocation is envy free
  Since agent 1 will pick the new bundle of agent 2 the potential increases again.

□

We must note that the same potential function appeared in [ACG+22].

5.3.3 EFX with $n + 2$ goods

The final result of this thesis is a proof that EFX allocations exist when the number of items is at most two more than the number of agents even when the agents have general valuations. As note earlier, Amanatidis et al., [AMN20], show this for additive valuations while Mahara, [Mah21], proved even for three more items but via a very long analysis.

Theorem 9. An EFX allocation exists when the number of goods is at most two higher than the number of agents.

Proof. We have discussed how an allocation with up to $n+1$ items is trivial so it remains to prove the statement for exactly $n+2$ items. Our proof is algorithmic and combines the allocation rule $\mathcal{U}$ presented in this paragraph with the idea of case analysis based on the number of sources. After allocating the first n items in a greedy style, and with a, b being the remaining two, there are 3 cases:

- Two sources in E_G
  Then we simply allocate one item to each.

- One source s in E_G and $\mathcal{U}(\mathcal{A}_s) \subset \mathcal{A}_s$
  Since there is a single source, the envious agent is reachable and we can allocate the last item to the new source after decycling the graph.

- One source s in E_G and $\mathcal{U}(\mathcal{A}_s) = \{a, b\}$
  Now we create a new allocation from scratch where the first agents up to the most envious pick as before, the envious agent gets a and b and the remaining agents pick again greedily, with the last one receiving the final item.

In all three cases, it easy to see that agents with one item prefer it over any item belonging to bundle of size 2, thus no strong envy exists. $\square$

Chapter 6

Experimental evaluation

The final chapter of the main body of this thesis is dedicated to some experiments. Looking back, every contribution presented in Chapter 5, with the exception of the simplifications of existing results in paragraph 5.3, is based upon some assumption. Therefore, we deem proper to examine how often our assumptions occur in real world data.

6.1 Setup

First of all, we should mention that the data used is extracted from the Fair Division website `www.spliddit.org`, [GP15]. We are thankful to Nisarg Shah for providing the data as of July 2021. The site provides a range of applications but, of course, we are focused on dividing goods. Now, our setup is pretty simple: since Spliddit allows for instances with both divisible and indivisible goods together we filtered out any instance that contains the first category and implemented our algorithms in Python, [VRD09]. We collected information about instances with distinct favorite goods (Theorem 3), bounded intervals ([ABFR$^+$21]) and about the tiers of the rankings (Theorem 7) which includes information about the top n ranking (Theorem 5).

6.2 About zero valued items

A quick pause before presenting the results. So far we have consider only goods, i.e., positive valued items, and have left negative ones out of our scope. But what about zero valued goods? The EFX definition of Caragiannis et al., [CKM$^+$16], considers only the hypothetical removal of a positive value good. Let us see it in practice.

Example 6.1. Consider the instance with 2 agents and 3 items

	a	b	c
Agent 1	3	2	1
Agent 2	0	0	1

According to the original definition, the allocation where agent 1 receives every item is EFX. Still, the allocation where agent 2 receives c seems fairer; and it is, indeed, envy free. △

Kyropoulou et al., [KSV20], strengthen the definition to include zero valued goods under the name EFX_0. Clearly, EFX_0 implies EFX but not vice versa. But how does zero valued items alternate our results?

To demonstrate their effect, let us consider an execution of pECE, algorithm 4, with input the instance of example 6.1. When there are more sources the algorithm picks one arbitrarily; therefore it is possible that agent 1 will be the first to receive one item and a will be her choice. Since $v_2(a) = 0$ the allocation remains envy free and it is possible that agent 1 is again the source. Continuing in the same manner we can see that the algorithm may output the allocation described in the example; which is $0-\text{EFX}_0$. Thus the bound of 1/2 does not hold anymore. To maintain consistency with the rest of the thesis, we have opted to perform our experiments with the original definition.

6.3 Results

At this point, we will remind the reader that EFX allocations are guaranteed to exist when there are few items: $m \leq n + 3$. Consequently, it would be expected to work only with the subset of "many items" instances. Unfortunately, their number is too small. As a middle solution, we decided to present the results firstly for the whole dataset and then for the many items instances only.

6.3.1 All instances

A similar discussion as the one above can be had about the number of agents since EFX allocations exist when there are at most three. In the same spirit as before, we will present the results for both cases. The total number of instances is 4323 and it drops down to 203 if we exclude the instances with few agents. The statistics are the following:

#agents	Distinct favorites	Bounded intervals	Same top n
n	1.851%	6.153%	13.046%
$n > 3$	0%	7.389%	3.941%

Table 6.1: Statistics for all instances

As for the information about the tiers, it is presented in Figure 6.1. For presentation's sake, we have chosen to include up to the tier of size 10. From the two diagrams we see that theorem 7 applies to roughly 29% of all instances while the number reduces to 18% when there are more than 3 agents.

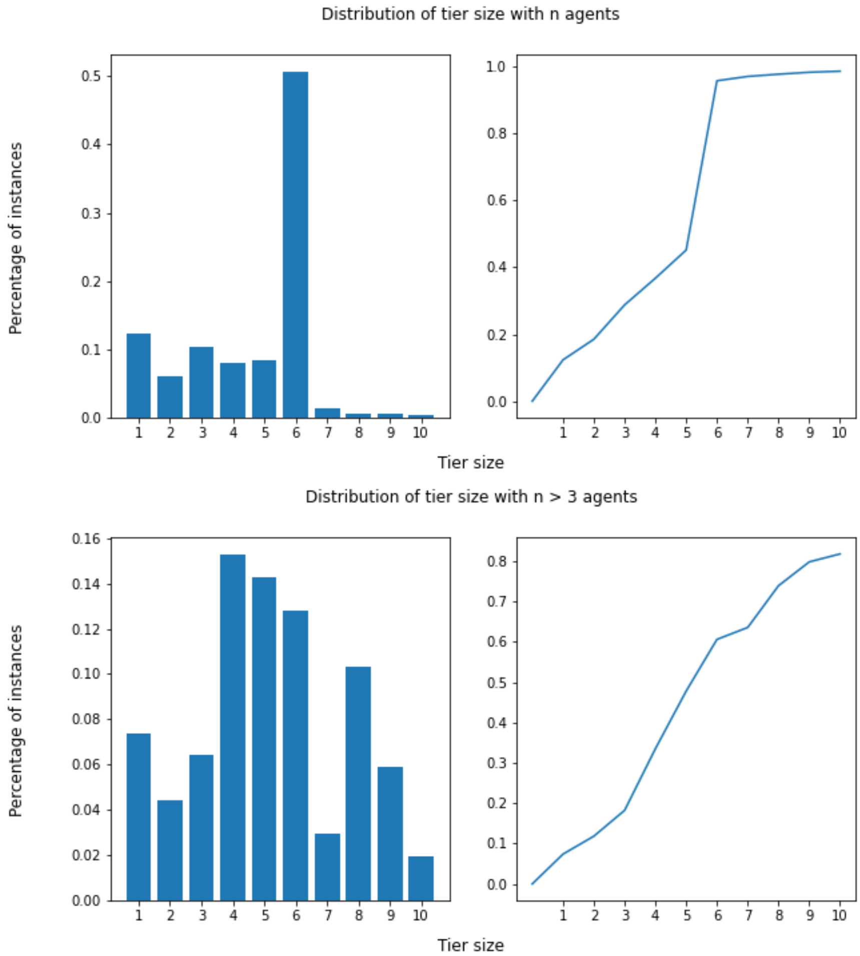

Figure 6.1: Tier information for all instances

6.3.2 Many items instances

Now we repeat our previous experiments excluding instances with few items ($m \leq n + 3$). There are 475 such instances and, unfortunately, only 51 if we further constraint them on the number of agents.

#agents	Distinct favorites	Bounded intervals	Same top n
n	3.158%	5.684%	12.842%
$n > 3$	0%	5.882%	3.922%

Table 6.2: Statistics for many items instances

In Figure 6.1 we see that $\approx 17.5\%$ of the many items instances have a tier size of at most 3. Moreover, only three of those many item instances include 4 or more agents. And, actually, in all three the agents share a common ranking. Therefore, algorithms 4 and 7 both output an EFX allocation for all 3 instances. Still, we would like to compare the actual approximation ratios versus the theoretical results for both algorithms so we had to

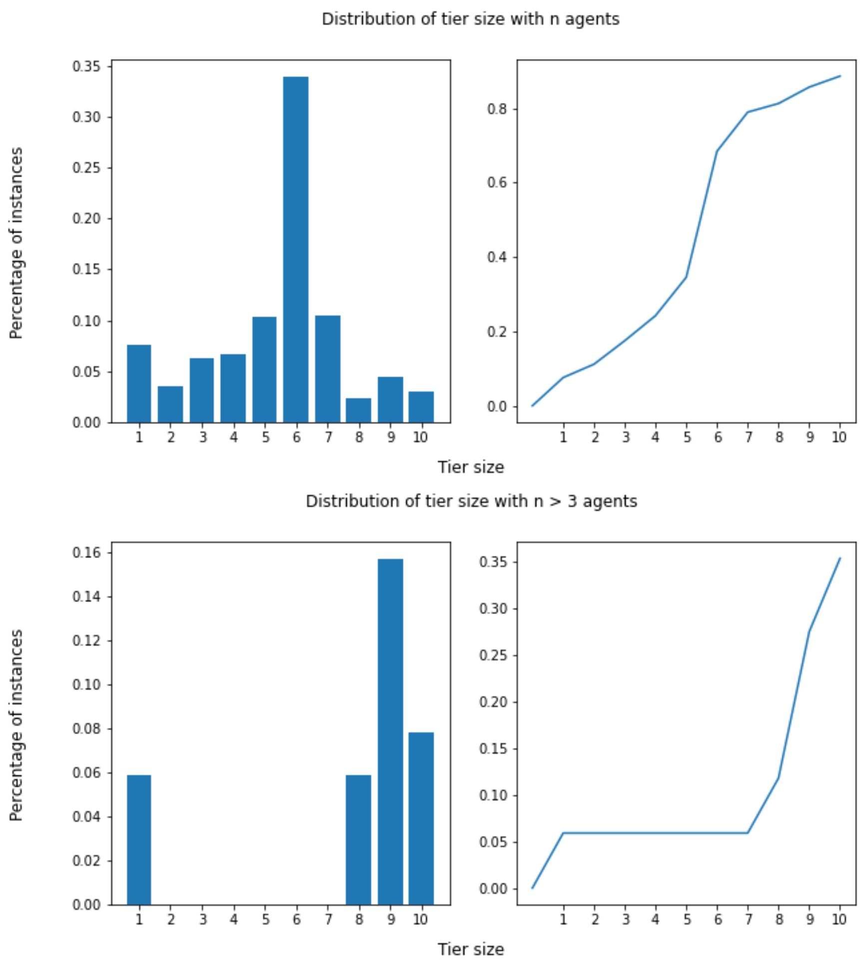

Figure 6.2: Tier information for many items instances

include instances with few agents. Interestingly enough, both algorithm ouptut an EFX allocation for 77 out of the 83 instances. The remaining 6 were split with an almost equal average approximation ratio. Finally, pECE computed an EFX allocation for 79 out of 83 instances and algorithm 7 for one more; with only one instance where neither method achieved a perfect ratio.

Chapter 7

Conclusion and Future Work

In this thesis, we studied the problem of fairly allocating a set of indivisible goods to agents with possibly different valuation functions. Out of the many fairness criteria we focused on the most prominent one: EFX. We showed a tight analysis of a variant of the famous Envy Cycle Elimination method, which eventually resulted in an approximation framework for additive valuations functions. Under our framework, we saw how the approximation version of EFX gradually degrades as some existence condition is getting more and more relaxed. Our main result was a $2/3$ approximation under the assumption that all n agents have a common ranking of the top n items. The main open question now is whether one can achieve the same result unconditionally.

Then, we moved on to the the broader class of cancelable valuation functions where we led a new path to attack the problem of the existence of EFX allocations. We introduced tiered rankings as a way to extend the setting with one identical ranking to families with exponentially many similar ones. Some future work could be the extension to tiers of size 4 and beyond. Ideally, it may be possible to show that that the tier size is as large as the number of agents; similarly to the top n ranking discussed above.

Finally, we also studied general valuation functions where we presented alternative proofs for a few existing results. The most interesting followup of this direction is applying our source based analysis to larger instances with few items, eg $m = n + 4$.

Bibliography

[ABFR+21] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. Maximum nash welfare and other stories about efx. *Theoretical Computer Science*, 863:69–85, 2021.

[ABFRV22] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A Voudouris. Fair division of indivisible goods: A survey. *arXiv preprint arXiv:2202.07551*, 2022.

[ABM18] Georgios Amanatidis, Georgios Birmpas, and Vangelis Markakis. Comparing approximate relaxations of envy-freeness. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 42–48. ijcai.org, 2018.

[ACG+22] Hannaneh Akrami, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. Efx allocations: Simplifications and improvements. *arXiv preprint arXiv:2205.07638*, 2022.

[ACIW22] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–21, 2022.

[ADG91] Ahmet Alkan, Gabrielle Demange, and David Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica: Journal of the Econometric Society*, pages 1023–1039, 1991.

[AK21] Noga Alon and Michael Krivelevich. Divisible subdivisions. *Journal of Graph Theory*, 98(4):623–629, 2021.

[AMN20] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. Multiple birds with one stone: Beating 1/2 for efx and gmms via envy cycle elimination. *Theoretical Computer Science*, 841:94–109, 2020.

[ARS22] Hannaneh Akrami, Rojin Rezvan, and Masoud Seddighin. An ef2x allocation protocol for restricted additive valuations. *arXiv preprint arXiv:2202.13676*, 2022.

[AW19] Martin Aleksandrov and Toby Walsh. Greedy algorithms for fair division of mixed manna. *arXiv preprint arXiv:1911.11005*, 2019.

[BBK22] Benjamin Aram Berendsohn, Simona Boyadzhiyska, and László Kozma. Fixed-point cycles and efx allocations. *arXiv preprint arXiv:2201.08753*, 2022.

[BBKN18] Siddharth Barman, Arpita Biswas, Sanath Krishnamurthy, and Yadati Narahari. Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[BCFF21] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. (almost full) efx exists for four agents (and beyond). *arXiv preprint arXiv:2102.10654*, 2021.

[BEF21] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5119–5126, 2021.

[BL16] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016.

[BT96] Steven John Brams and Alan D Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.

[Bud11] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[CCLW19] Hau Chan, Jing Chen, Bo Li, and Xiaowei Wu. Maximin-aware allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, IJCAI'19, page 137–143. AAAI Press, 2019.

[CGH19] Ioannis Caragiannis, Nick Gravin, and Xin Huang. Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 527–545, 2019.

[CGM20] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. Efx exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 1–19, 2020.

[CGM+21] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. Improving efx guarantees through rainbow cycle number. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 310–311, 2021.

[CKM+16] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322, 2016.

[CKMS21] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4):1336–1358, 2021.

[GM21] Jugal Garg and Aniket Murhekar. Computing fair and efficient allocations with few utility values. In *International Symposium on Algorithmic Game Theory*, pages 345–359. Springer, 2021.

[GMT14] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. Near fairness in matroids. In *ECAI*, pages 393–398, 2014.

[GP15] Jonathan Goldman and Ariel D Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2015.

[GS22] Jugal Garg and Eklavya Sharma. Existence and computation of epistemic efx. *arXiv preprint arXiv:2206.01710*, 2022.

[HSVX21] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Fair and efficient allocations under lexicographic preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5472–5480, 2021.

[KSV20] Maria Kyropoulou, Warut Suksompong, and Alexandros A Voudouris. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841:110–123, 2020.

[LMMS04] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.

[Mah20] Ryoga Mahara. Existence of efx for two additive valuations. *arXiv preprint arXiv:2008.08798*, 2020.

[Mah21] Ryoga Mahara. Extension of additive valuations to general valuations on the existence of efx. In *29th Annual European Symposium on Algorithms (ESA 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[Mar17] Evangelos Markakis. Approximation algorithms and hardness results for fair division with indivisible goods. *Trends in Computational Social Choice*, pages 231–247, 2017.

[Mou04] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.

[MS21] Tamás Mészáros and Raphael Steiner. Zero sum cycles in complete digraphs. *European Journal of Combinatorics*, 98:103399, 2021.

[PR18] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2584–2603, 2018.

[Pro15] Ariel D Procaccia. Cake cutting algorithms. In *Handbook of computational social choice, chapter 13*. Citeseer, 2015.

[Pro20] Ariel D Procaccia. Technical perspective: An answer to fair division's most enigmatic question. *Communications of the ACM*, 63(4):118–118, 2020.

[RW98] Jack Robertson and William Webb. *Cake-cutting algorithms: Be fair if you can*. AK Peters/CRC Press, 1998.

[Ste48] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

[Vaz01] Vijay V Vazirani. *Approximation algorithms*, volume 1. Springer, 2001.

[VRD09] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.