

Tutorial on multiple currency accounting

Author: Peter Selinger <selinger at users dot sourceforge dot net>

Written: Jun 4, 2005

Updated: Oct 20–26, 2008 (added section on realized and unrealized gains)

Updated: Nov 15, 2010 (changed sectioning)

Updated: Jan 2, 2011 (added Australia info to Section 5.3)

Updated: Jun 2, 2015 (fixed typo in Table 5.2, reported by Alex Chiang)

Updated: Aug 14, 2015 (fixed typo in Table 3.2, reported by Francisco Milan Campos)

Contents

0. Introduction

1. A brief review of double-entry accounting

- 1.1. [Single-entry accounting](#)
- 1.2. [Net worth and equity](#)
- 1.3. [Double entry accounting](#)
- 1.4. [Income and Expenses](#)
- 1.5. [The Fundamental Accounting Equation](#)

2. Foreign currency accounting issues

- 2.1. [Single-entry foreign currency accounting](#)
- 2.2. [Double entry foreign currency accounting, the wrong way](#)
- 2.3. [Foreign exchange gains and losses](#)
- 2.4. [Calculating gains and losses through adjustments](#)

3. Foreign currency accounting according to SSAP 20

- 3.1. [Accounting for non-monetary current assets](#)
- 3.2. [The "official" solution: translation to a single currency](#)

4. Foreign currency accounting using currency trading accounts

- 4.1. [Restoring symmetry to foreign currency accounting](#)
- 4.2. [Accounting with currency trading accounts](#)
- 4.3. [Equivalence with SSAP 20](#)
- 4.4. [Example: tracking exchange gains and losses by source](#)

5. Realized and unrealized gains and losses

- 5.1. [An example from real estate](#)
- 5.2. [Importance of the "identity" of assets](#)
- 5.3. [The FIFO, LIFO, and adjusted cost base methods](#)

6. Translation issues

- 6.1. [Independence of viewpoint](#)
- 6.2. [Reference currency vs. Equity currency](#)

7. Specific software

- 7.1. [GnuCash](#)

0. Introduction

Accounting with multiple currencies can be tricky. It is easy enough to use an exchange rate to translate between different currencies. But things get less obvious when one takes into account how exchange rates fluctuate over time. Multi-currency transactions can incur foreign exchange gains or losses. How does one account for these, while respecting the Fundamental Accounting Equation? Do different accounting methods

give different results? What are their advantages and disadvantages? Does one need to pick a single reference currency, and why? And most importantly, how does all of this work?

As a non-expert, I found it surprisingly difficult to find useful information on this subject. I searched online, and I also consulted about fifteen different accounting textbooks in the U.K. and Canada. Of the few books that mentioned foreign currencies at all, most contained only a single paragraph on the topic, stating that foreign currency transactions should be translated to the local currency. I also studied the foreign currency support in the accounting software I was using (GnuCash), and I quickly found that the software did not balance my accounts correctly when multiple currencies were involved.

After much frustration and some additional thought and research, I believe I arrived at an understanding how multiple currency accounting should work, and why. So I wrote this document as a resource for those who might be looking for such information. The information in this tutorial is a synthesis of what I learned from books, what I learned from accounting standards such as SSAP 20, and what I figured out by myself. Since I am not an accounting expert, I am sure that some incorrect terminology has slipped into this document. If you find a mistake, I would be happy if you send me a correction. Aside from terminology issues, I am reasonably certain that the concepts described in this tutorial are sound.

Disclaimer: I am not an accountant, and I am not qualified to offer any kind of accounting advice. This tutorial expresses my opinion on how accounting with multiple currencies should work. I do not claim anything about the accuracy or legality of any of the information presented here, nor do I warrant its correctness or applicability in any way. Although I cite some accounting standards documents, I do not claim that this tutorial conforms with accepted accounting standards. If you need accounting advice, please consult a professional accountant.

1. A brief review of double-entry accounting

1.1. Single-entry accounting

Users of accounting software fall into two main classes: home users and business users. Most home users are only interested in tracking *assets* and *liabilities*. Assets are what you *own* (such as money in your bank account or cash in your pocket), and liabilities are what you *owe* (such as the balance on your credit card account, or money you owe your neighbor). An *account* keeps a running total of one particular type of asset or liability. A *transaction* is an increase or decrease of the balance of one or more accounts.

Table 1.1: Single-entry personal accounting

Date	Description	Bank Account (asset)	Cash (asset)	Credit Card (liability)
Jan 1	Opening balance	\$1,000	\$20	\$600
Jan 4	Get paid	+ \$200		
	Balance	\$1,200	\$20	\$600
Jan 5	Buy food with credit card			+ \$70
	Balance	\$1,200	\$20	\$670
Jan 12	Withdraw cash	- \$100	+ \$100	
	Balance	\$1,100	\$120	\$670
Jan 15	Pay credit card bill	- \$670		- \$670
	Balance	\$430	\$120	\$0
Jan 18	Get paid	+ \$200		
	Balance	\$630	\$120	\$0
Jan 20	Buy food with cash		- \$65	
	Balance	\$630	\$55	\$0
Jan 23	Buy book with cash		- \$16	
	Balance	\$630	\$39	\$0

[Table 1.1](#) shows an example of single-entry personal accounting with three accounts. It uses just a single currency, the Canadian Dollar. Note that account balances are shown on yellow lines, and transactions are shown on white lines. Also note that in asset accounts, positive numbers mean "more assets", whereas in liability accounts, positive numbers mean "more liability". Thus, transferring money from one asset account to another (as on January 12) leads to an increase in one account and a decrease in the other. But transferring money from an asset account to a liability account (as in paying a credit card bill on January 15) leads to a decrease in both accounts.

1.2. Net worth and equity

A person's *net worth* is equal to their total assets minus their total liabilities. In other words, the net worth is the effective amount of money (and property) they own after compensating for all outstanding debts. A person's net worth can be negative if the total liabilities exceed the total assets (this is called "living above one's means", and is not recommended). In a business context, net worth is also called *equity*. We therefore have the first accounting equation:

$$\text{Assets} - \text{Liabilities} = \text{Equity}.$$

Going back to the example from [Table 1.1](#), note that some transactions, like "get paid" on January 4 and 18, increase a person's net worth. This is called *income*. Other transactions, like "buy food" or "buy books" on January 5, 20, and 23, decrease a person's net worth. This is called an *expense*. Still other transactions, such as "withdraw cash" on January 12 or "pay credit card bill" on January 15, do not affect the net worth at all.

The following chart shows the same transactions as [Table 1.1](#), but this time we also show the net worth in a separate column.

Table 1.2: Personal accounting with net worth shown

Date	Description	Bank Account (asset)	Cash (asset)	Credit Card (liability)	Net Worth (equity)
Jan 1	Opening balance	\$1,000	\$20	\$600	\$420
Jan 4	Get paid	+ \$200			+ \$200
	Balance	\$1,200	\$20	\$600	\$620
Jan 5	Buy food with credit card			+ \$70	- \$70
	Balance	\$1,200	\$20	\$670	\$550
Jan 12	Withdraw cash	- \$100	+ \$100		
	Balance	\$1,100	\$120	\$670	\$550
Jan 15	Pay credit card bill	- \$670		- \$670	
	Balance	\$430	\$120	\$0	\$550
Jan 18	Get paid	+ \$200			+ \$200
	Balance	\$630	\$120	\$0	\$750
Jan 20	Buy food with cash		- \$65		- \$65
	Balance	\$630	\$55	\$0	\$685
Jan 23	Buy book with cash		- \$16		- \$16
	Balance	\$630	\$39	\$0	\$669

Note that the accounting equation "Assets – Liabilities = Equity" is not only satisfied by the balances (yellow lines), but also by the individual transactions (white lines).

1.3. Double entry accounting

Let us promote the equity column from [Table 1.2](#) by calling it an "account". Then [Table 1.2](#) is a simple instance of *double entry accounting*. The idea is to keep a separate running total for equity, in addition to the running totals for assets and liabilities. Each transaction must be *balanced*, which means, it must satisfy the accounting equation.

Historically (i.e., before computers were used), one of the purposes of double entry accounting was to catch arithmetic errors and data entry errors. Since each transaction had to be recorded in at least two places, the results of additions and subtractions could be double-checked simply by checking whether the accounts were balanced.

Today, arithmetic mistakes are no longer a major issue (although data entry errors still are). However, there is one other important benefit of double-entry accounting. Namely, it gives a convenient way of tracking and classifying income and expenses.

1.4. Income and Expenses

For businesses, it is important to keep track not just of assets and liabilities, but also of income and expenses. For example, a business needs to classify its income according to source. This is necessary both for tax reasons (income from different sources may be taxed differently), and also for business reasons (it helps to determine which parts of the business are profitable).

In the asset/liability method of tracking finances, money is classified according to its *location*: how much money is in your bank account, and how much in your wallet? By contrast, in tracking income and expenses, one must classify money according to its *purpose*: how much money did you spend on food, and how much on books?

For example, in [Table 1.2](#), the total money spent on food in the given time period was \$135 (namely, \$70 on January 5 and \$65 on January 20). Some of the food was bought with cash, and some with a credit card. While these different payment methods matter from the point of view of assets and liabilities, they are irrelevant from the point of view of expenses.

The most convenient way to track income and expenses is to split the column for "equity" from [Table 1.2](#) into several columns. This is normally done for a fixed time period, for example, one year (called an "accounting period"). The equity at the beginning of the time period is called the "initial capital". All increases in equity since the beginning of the time period are called "income", and all decreases are called "expenses". We therefore have the second accounting equation:

$$\text{Equity} = \text{Capital} + \text{Income} - \text{Expenses}.$$

We will therefore split the "equity" account into three columns called "initial capital", "income" and "expense". The income and expense columns can be further subdivided to account for different types of income and expenses. To keep our example simple, we will consider a single income account called "salary", and two expense accounts called "food" and "books".

Table 1.3: Personal accounting with income/expense tracking

Date	Description	Bank Account (asset)	Cash (asset)	Credit Card (liability)	Initial Capital (capital)	Salary (income)	Food (expense)	Books (expense)
Jan 1	Opening balance	\$1,000	\$20	\$600	\$420	\$0	\$0	\$0
Jan 4	Get paid	+ \$200				+ \$200		
	Balance	\$1,200	\$20	\$600	\$420	\$200	\$0	\$0
Jan 5	Buy food with credit card			+ \$70			+ \$70	
	Balance	\$1,200	\$20	\$670	\$420	\$200	\$70	\$0
Jan 12	Withdraw cash	- \$100	+ \$100					
	Balance	\$1,100	\$120	\$670	\$420	\$200	\$70	\$0
Jan 15	Pay credit card bill	- \$670		- \$670				

	Balance	\$430	\$120	\$0	\$420	\$200	\$70	\$0
Jan 18	Get paid	+ \$200				+ \$200		
	Balance	\$630	\$120	\$0	\$420	\$400	\$70	\$0
Jan 20	Buy food with cash		– \$65				+ \$65	
	Balance	\$630	\$55	\$0	\$420	\$400	\$135	\$0
Jan 23	Buy book with cash		– \$16					+ \$16
	Balance	\$630	\$39	\$0	\$420	\$400	\$135	\$16

As there is no limit to the number of accounts one can have, there is no limit on the number of different types of income and expenses that can be tracked using this system. Note that expenses (like liabilities) are recorded as positive numbers, even though they contribute to equity negatively.

Also note that the "initial capital" account never changes in this example. In business use, there are other types of capital besides "initial capital", and some of these other types are changeable. For example, if somebody invests money in a business (thereby becoming an owner), then this is considered an increase in capital rather than income. Thus, in general, there can be more than one capital account.

1.5. The Fundamental Accounting Equation

Putting together the first and second accounting equations

$$\text{Assets} - \text{Liabilities} = \text{Equity}$$

and

$$\text{Equity} = \text{Capital} + \text{Income} - \text{Expenses},$$

we obtain the following, which is called the Fundamental Accounting Equation:

$$\text{Assets} - \text{Liabilities} = \text{Capital} + \text{Income} - \text{Expenses}.$$

The reader should check that both the yellow and the white rows in [Table 1.3](#) indeed satisfy this equation.

Remark: "equity" is sometimes called "capital" and vice versa; I am not really sure of the most correct terminology. In this tutorial, I will continue to use these terms with their meanings introduced above.

Remark: Traditionally, accountants did not like to subtract or to use negative numbers. The accounting equation is therefore often written as "Assets + Expenses = Capital + Income + Liabilities". The quantities on the left-hand side of this last equation are sometimes called *activa* (Assets and Expenses), and those on the right-hand side are sometimes called *passiva* (Capital, Income, and Liabilities). The reader will easily verify that the two versions of the accounting equation are mathematically equivalent. Conceptually and pedagogically, I find the division into *location* (Assets and Liabilities) and *purpose* (Capital, Income, and Expenses) more useful than the division into *activa* and *passiva*. Therefore, in this tutorial, I will continue to use the version of the accounting equation that uses subtraction.

2. Foreign currency accounting issues

2.1. Single-entry foreign currency accounting

Handling foreign currencies in single-entry accounting is very simple. Each asset or liability account can have its own currency. When transferring money from one account to another, we simply use the exchange rate that is in effect at the time of the transfer.

Suppose you keep two kinds of cash in your pocket: Canadian dollars (CAD) and U.S. dollars (USD). On January 1, you have CAD 200 and no U.S. dollars. On January 2, you exchange CAD 120 for USD 100 at a bank (at the day's exchange rate, which is 1.20 CAD/USD). On January 3, you spend USD 40 on food. On January 5, you exchange USD 60 for CAD 75 (at that day's rate of 1.25 CAD/USD). On January 7, you spend CAD 20 on food. These transactions can easily be accounted for with two asset accounts:

Table 2.1: Single-entry foreign currency accounting

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)
Jan 1	Opening balance	CAD 200	USD 0
Jan 2	Exchange (1 USD = 1.20 CAD)	- CAD 120	+ USD 100
	Balance	CAD 80	USD 100
Jan 3	Buy food		- USD 40
	Balance	CAD 80	USD 60
Jan 5	Exchange (1 USD = 1.25 CAD)	+ CAD 75	- USD 60
	Balance	CAD 155	USD 0
Jan 7	Buy food	- CAD 20	
	Balance	CAD 135	USD 0

As the example shows, foreign currencies are easy to account for as long as you only want to keep track of *assets* (and it is equally easy for *liabilities*). The process of doing currency transfers with exchange rates can easily be automated in software.

2.2. Double entry foreign currency accounting, the wrong way

Things get more complicated if we also want to keep track of income and expenses. To illustrate this, we follow the method from [Table 1.3](#) and add accounts for initial capital and expenses. (Since there is no income in this example, we don't need an income account). Because I live in Canada, we will keep the capital and expense accounts in Canadian dollars. Of course, when there is a transaction that affects one of these accounts, the amount will have to be translated to Canadian dollars. To do this, we have to know the exchange rate on January 3; let us assume for the sake of this example that this exchange rate was 1.30 CAD/USD, so that USD 40 = CAD 52 on that day.

Table 2.2: Double-entry foreign currency accounting, naively

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Food (expense)
Jan 1	Opening balance	CAD 200	USD 0	CAD 200	CAD 0
Jan 2	Exchange (1 USD = 1.20 CAD)	- CAD 120	+ USD 100		
	Balance	CAD 80	USD 100	CAD 200	CAD 0
Jan 3	Buy food (1 USD = 1.30 CAD)		- USD 40		+ CAD 52
	Balance	CAD 80	USD 60	CAD 200	CAD 52
Jan 5	Exchange (1 USD = 1.25 CAD)	+ CAD 75	- USD 60		
	Balance	CAD 155	USD 0	CAD 200	CAD 52
Jan 7	Buy food	- CAD 20			+ CAD 20
	Balance	CAD 135	USD 0	CAD 200	CAD 72

Note that we have recorded each transaction according to the exchange rate in effect on the date the transaction was made. Remember the accounting equation "Assets - Liabilities = Capital + Income - Expenses". Please convince yourself that each transaction (white line) indeed satisfies this equation, if one takes the daily exchange rate into account. Therefore, the white lines are balanced.

So what is the problem? The problem is that the yellow lines are not balanced. For example, consider the balance for January 7. We started with CAD 200, and we spent CAD 72 on food, so we should still have

CAD 128 in our pocket. Yet, we have CAD 135. So something went wrong: we gained CAD 7 without any record of it on the "equity" side of the balance sheet. The accounting system of [Table 2.2](#) violates the Fundamental Accounting Equation.

2.3. Foreign exchange gains and losses

The CAD 7 that were "won" in [Table 2.2](#) are called a *foreign exchange gain*. Such gains are caused by exchange rate fluctuations between transactions. For example, if you buy a foreign currency, then sell it later at a higher price, you have enjoyed a foreign exchange gain. Similarly, if you lose money due to exchange rate fluctuations, then this is called a *foreign exchange loss*. For accounting purposes, foreign exchange gains count as *income*, and foreign exchange losses count as an *expense*. What the accounting method from [Table 2.2](#) lacks is a way to track foreign exchange gains and losses.

The question we need to address is therefore how to calculate foreign exchange gains and losses, and how to record them, preferably while keeping the Fundamental Accounting Equation intact.

Calculating foreign exchange gains and losses is somewhat tricky, because such gains and losses do not occur during transactions, but rather in the time intervals *between* transactions. It is therefore not immediately clear at what time they should be entered in the balance sheet. In principle, gains and losses happen every time the exchange rate changes, whether or not there was a transaction. But it is clearly not feasible to record them on a continuous basis (for example, once per second). So there must be a better way.

Another potential complication is that the concept of a gain or loss depends on the "base" currency, and therefore on your point of view. Suppose you have CAD 150 and USD 100, and the exchange rate changes from 1.20 CAD/USD to 1.25 CAD/USD. If you are Canadian, then you have just enjoyed a gain (your equity increased from CAD 270 to CAD 275). But if you are American, then you have suffered a loss (your equity decreased from USD 225 to USD 220). This dependency on a base currency seems necessary, but it is also frustrating, as it seems to imply that a multi-currency accounting system cannot be "symmetric" in the way different currencies are treated. (It is in fact possible to design a symmetric system, as we will see in [Section 6.1](#) below).

2.4. Calculating gains and losses through adjustments

In principle, one could account for foreign exchange gains and losses in the same way one accounts for intractable arithmetic errors: by adding an "adjusting entry" before each yellow line in [Table 2.2](#), forcing the balances to obey the accounting equation. This method is shown in [Table 2.3](#).

Table 2.3: Foreign currency accounting with adjustments

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Food (expense)	Exchange gain (income)
Jan 1	Opening balance	CAD 200	USD 0	CAD 200	CAD 0	CAD 0
Jan 2	Exchange (1 USD = 1.20 CAD)	- CAD 120	+ USD 100			
	Balance	CAD 80	USD 100	CAD 200	CAD 0	CAD 0
Jan 3	Buy food (1 USD = 1.30 CAD)		- USD 40		+ CAD 52	
	Adjustment					+ CAD 10
	Balance	CAD 80	USD 60	CAD 200	CAD 52	CAD 10
Jan 5	Exchange (1 USD = 1.25 CAD)	+ CAD 75	- USD 60			
	Adjustment					- CAD 3
	Balance	CAD 155	USD 0	CAD 200	CAD 52	CAD 7
Jan 7	Buy food	- CAD 20			+ CAD 20	
	Balance	CAD 135	USD 0	CAD 200	CAD 72	CAD 7

To calculate [Table 2.3](#), we simply had to look at each day's balance (yellow line), and force it to satisfy the accounting equation according to that day's exchange rate. An adjusting entry was then added to produce the desired numbers. As a result, the "yellow" (balance) lines are balanced, but the "white" (transaction) lines are not.

Note that there was an exchange loss of CAD 3 on January 5. While this technically is an "expense", we have not created a special "expense" account for it; instead we have recorded this loss as a "negative gain". This is in accordance with an accepted accounting principle that states that exchange gains and losses can be directly offset against each other.

The method shown in [Table 2.3](#) is undesirable for several reasons. The main problem is that the additional "adjusting entries" are themselves not balanced, thus violating the principle of double entry accounting. This removes a source of redundancy. In effect, since one constantly has to "force" the accounts to balance, one can no longer notice the difference between currency fluctuations and random arithmetical errors.

Moreover, even if the accounts are in balance at one point, they continually "go out of balance" as the exchange rate changes, even while no transactions occur. Thus, adjustment entries should really be added continuously, rather than once per transaction. This would be an undesirable thing to have to do.

Another problem with the adjustment method is that it makes it difficult to add or change transactions retroactively. Suppose we want to enter another transaction on January 2. Then all future adjusting balances (on January 3 and 5) have to be re-calculated. While it would be possible for a computer to do this automatically, it is not very elegant and even seems manipulative. Instead of simply *recording* transactions, we are now *calculating* transactions, and the separation of data entry and analysis seems to be violated. Cause and effect of individual transactions are difficult to track.

3. Foreign currency accounting according to SSAP 20

3.1. Accounting for non-monetary current assets

So if the methods shown in [Tables 2.2](#) and [2.3](#) are bad, then how should double-entry accounting for multiple currencies be done? Before we look at possible solutions, it is useful to recall how one accounts for the changing value of non-monetary assets.

Imagine that you are in the business of buying and selling oil. Let us assume that oil is a *current* asset, which means that it can be easily converted to cash on the market. Therefore, oil should always be accounted for at its *current market price*. (This is different from *non-current* assets, such as real-estate, which are usually accounted for at their historical cost).

Suppose you buy oil at market prices, and you sell it at market price plus a profit margin of CAD 5 per barrel. On January 1, you buy 10 barrels of oil at a price of CAD 30 per barrel. On January 2, you sell 2 barrels for CAD 35 each (the market price is still CAD 30). On January 3, the price of oil increases to CAD 32 per barrel. On January 4, you sell 5 more barrels at CAD 37. We can account for this as follows:

Table 3.1: Accounting for non-monetary current assets

Date	Description	Cash (asset)	Oil (asset)	Inventory	Initial capital (capital)	Sales profit (income)	Capital gain (income)
Jan 1	Opening balance	CAD 500	CAD 0	0 barrels	CAD 500	CAD 0	CAD 0
Jan 1	Buy 10 barrels	- CAD 300	+ CAD 300	+10 barrels (@CAD 30)			
	Balance	CAD 200	CAD 300	10 barrels (@CAD 30)	CAD 500	CAD 0	CAD 0
Jan 2	Sell 2 barrels	+ CAD 70	- CAD 60	-2 barrels (@CAD 30)		+ CAD 10	
				8 barrels			

	Balance	CAD 270	CAD 240	(@CAD 30)	CAD 500	CAD 10	CAD 0
Jan 3	Re-value inventory		+ CAD 16	8 barrels (+CAD 2)			+ CAD 16
	Balance	CAD 270	CAD 256	8 barrels (@CAD 32)	CAD 500	CAD 10	CAD 16
Jan 4	Sell 5 barrels	+ CAD 185	- CAD 160	-5 barrels (@CAD 32)		+ CAD 25	
	Balance	CAD 455	CAD 96	3 barrels (@CAD 32)	CAD 500	CAD 35	CAD 16

There are several things to note in this example. The first remark is that the green column marked "Inventory" is not an *account*; it is simply an informational column recording the number of barrels of oil we currently own, and their current price. This column does not usually appear in a company's accounts; it can be kept separately. The accounts only contain the *value* of the assets, rather than the assets themselves.

Also note that this accounting method clearly distinguishes different sources of income: CAD 35 of the income is from sales profits (for 7 barrels sold at a profit margin of CAD 5 each), and CAD 16 is from capital gains (because the price of oil rose by CAD 2 at a time when you had 8 barrels).

From the viewpoint of multiple currency accounting, we remark that a "barrel of oil", while cumbersome to carry in your wallet, might nevertheless be regarded as a kind of currency. We make the following observations:

- All accounts are in CAD. Even though there is another currency involved, namely "barrels of oil", we do not actually keep any account in this currency. Our "oil asset" account is kept in CAD.
- Because any fluctuations in the price of oil are recorded in CAD, they are reflected on the asset side of the balance sheet and also on the income/expense side. As a result, the accounts are balanced at all times.
- Since the oil asset account is denominated in CAD and not in barrels, we must separately keep an inventory of the actual amount of oil owned, and also of the current market price. This inventory, which is shown in the green column in the example, is only informational; it is not part of the account sheet, and does not take part in any balancing.

3.2. The "official" solution: translation to a single currency

Most countries have a code of standard accounting practices that deals with foreign currency accounting. These can be searched for (and in some cases, found) on the web. Examples of such documents are SSAP 20 (for Canada and the U.K.) and FASB 52 (for the U.S.). It is my understanding that these various country-specific documents more or less all say the same thing. For example, SSAP 20 states:

During an accounting period, a company may enter into transactions which are denominated in a foreign currency. The result of each transaction should normally be translated into the company's local currency using the exchange rate in operation on the date on which the transaction occurred [...].

In effect, this means that foreign currencies should be treated precisely in the same way as non-monetary assets, i.e., they should be accounted for *in the local currency*. (As before, we use Canadian dollars as the local currency in our examples). This means that, just as the oil account in [Table 3.1](#) was denominated in CAD (and not in "barrels of oil"), the U.S. dollar account in our foreign currency example should also be denominated in CAD (and not in USD), when using this accounting method. Currency gains and losses are determined in the same way as other kinds of capital gains and losses: namely, by re-valuing accounts when a change of exchange rates occurs.

We illustrate this by reconsidering the example from [Tables 2.1–2.3](#), using the SSAP 20 method. Note that the accounting principle used here is very similar to [Table 3.1](#).

Table 3.2: Foreign currency accounting à la SSAP 20

Date	Description	CAD Cash (asset)	USD Cash (asset)	USD Cash Inventory	Initial Capital (capital)	Food (expense)	Exchange gain (income)
Jan 1	Opening balance	CAD 200	CAD 0	USD 0	CAD 200	CAD 0	CAD 0
Jan 2	Exchange	- CAD 120	+ CAD 120	+ USD 100 (@CAD 1.20)			
	Balance	CAD 80	CAD 120	USD 100 (@CAD 1.20)	CAD 200	CAD 0	CAD 0
Jan 3	Re-value USD		+ CAD 10	USD 100 (+ CAD 0.10)			+ CAD 10
	Balance	CAD 80	CAD 130	USD 100 (@CAD 1.30)	CAD 200	CAD 0	CAD 10
	Buy food		- CAD 52	- USD 40 (@CAD 1.30)		+ CAD 52	
	Balance	CAD 80	CAD 78	USD 60 (@CAD 1.30)	CAD 200	CAD 52	CAD 10
Jan 5	Re-value USD		- CAD 3	USD 60 (- CAD 0.05)			- CAD 3
	Balance	CAD 80	CAD 75	USD 60 (@CAD 1.25)	CAD 200	CAD 52	CAD 7
	Exchange	+ CAD 75	- CAD 75	- USD 60 (@CAD 1.25)			
	Balance	CAD 155	CAD 0	USD 0	CAD 200	CAD 52	CAD 7
Jan 7	Buy food	- CAD 20				+ CAD 20	
	Balance	CAD 135	CAD 0	USD 0	CAD 200	CAD 72	CAD 7

Compare this carefully to [Table 2.3](#). The first thing to notice is that each white and yellow row is balanced in [Table 3.2](#), which was not the case in [Table 2.3](#). On the equity side of the balance sheet (to the right of the solid black line), nothing has changed (except the order of some transactions). On the assets/liabilities side (to the left of the solid black line), the USD account is now denominated in CAD, and therefore, the currency fluctuations on January 3 and 5 lead to re-valuations of assets. As in [Table 3.1](#), the green column is only informational and should be disregarded for balancing purposes.

We summarize the main points of the SSAP 20 method:

- There is always one special *reference currency*, typically "the company's local currency".
- All accounts are kept in the reference currency.
- Tallies in non-reference currencies are only informational, and do not take part in balancing.
- Fluctuations in the relative values of non-reference currencies are reflected immediately in their corresponding asset accounts (and also in a currency gain/loss account).

The advantages of this method are:

- Since all entries are in one currency, no special software support for multiple currencies is needed — in principle, only single-currency accounting software is required.
- The accounts are balanced.

The disadvantages are:

- Separate "tallies" of foreign currencies and their exchange rates must be kept (the green columns in the examples). In fact, these "tallies" contain original (*source*) information, whereas the corresponding asset/liability accounts contain derived (*calculated*) information. In a sense, the foreign currency accounts that are kept in a local currency are only "virtual" accounts, whose contents are entirely calculated.

- Accounts must be re-valued periodically, and the question arises when to do this. In principle, this should happen just before every transaction, and just before generating a report. If there are many foreign currency accounts, this can mean a lot of (manual) work.
- It is relatively difficult to add or change a transaction retroactively. If new a transaction is entered for an earlier date, then all future currency gains and losses must be recalculated. This changes future transactions, not just future balances.
- There is an unsatisfying asymmetry in this accounting method, because one currency must be given prominence above all others.

The first three of these disadvantages can be alleviated by using software that supports the SSAP 20 accounting method. In this case, the user only has to enter the data from the green columns, and the software is able to take care of conversions, re-valuations, and calculating the corresponding gains and losses.

4. Foreign currency accounting using currency trading accounts

4.1. Restoring symmetry to foreign currency accounting

In [Section 3](#), one of our main complaints about the SSAP 20 accounting method was that it introduces an apparent asymmetry by requiring one currency to be singled out as the "reference currency". We would prefer to work with accounts that are directly denominated in multiple currencies, but we have seen in [Section 2](#) that this leads to problems with the accounting equation. Let us now investigate a method that restores symmetry while also respecting the accounting equation.

Consider the following simple scenario. On January 1, you have assets of CAD 60 and USD 100, and the exchange rate is 1.20 CAD/USD. Therefore, your total capital, in Canadian dollars, is CAD 60 + CAD 120 = CAD 180. Let us assume that you do not make any further transactions in January, but the exchange rate changes to 1.30 CAD/USD, 1.25 CAD/USD, and 1.15 CAD/USD on the three days following January 1. Recording only balances (yellow lines) and ignoring transactions (white lines) for the moment, we find that we should record the following, more or less in the style of [Table 2.3](#):

Table 4.1: Cumulative exchange gain as a function of exchange rate

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Exchange gain (income)
Jan 1	Balance (1 USD = 1.20 CAD)	CAD 60	USD 100	CAD 180	CAD 0
Jan 2	Balance (1 USD = 1.30 CAD)	CAD 60	USD 100	CAD 180	CAD 10
Jan 3	Balance (1 USD = 1.25 CAD)	CAD 60	USD 100	CAD 180	CAD 5
Jan 4	Balance (1 USD = 1.15 CAD)	CAD 60	USD 100	CAD 180	– CAD 5

Note that the only change from day to day is the exchange gain; there are no changes in the asset accounts or in the initial capital account. Also note that the accounts are balanced on each day, if one takes that day's exchange rate into account.

Now let us save some work by giving a *formula* for the exchange gain, instead of calculating it explicitly each day. Let the exchange rate on a given day be x CAD/USD, and let y be the corresponding balance in the exchange gain account. What is the formula for y in terms of x ? We must satisfy the Fundamental Accounting Equation "Assets – Liabilities = Capital + Income – Expense", and therefore we must have:

$$\text{CAD } 60 + x * \text{USD } 100 = \text{CAD } 180 + y.$$

This can easily be solved for y :

$$y = x * \text{USD } 100 - \text{CAD } 120.$$

In other words, the balance y of the exchange gain account on any given day is exactly equal to the current value of (USD 100 minus CAD 120), expressed in Canadian dollars. We can simplify [Table 4.1](#) by

expressing the exchange gain as a difference of multiple currencies, rather than expressing it as a value in a single currency.

Table 4.2: Cumulative exchange gain expressed as a difference of multiple currencies

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Exchange gain (income)
Jan 1	Balance (1 USD = 1.20 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 2	Balance (1 USD = 1.30 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 3	Balance (1 USD = 1.25 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 4	Balance (1 USD = 1.15 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120

Note that the exchange gains shown in [Table 4.2](#) are exactly identical to those shown in [Table 4.1](#), given each day's exchange rate. For example, on January 3, [Table 4.2](#) shows (USD 100 – CAD 120), but this is equal to the amount of CAD 5 shown in [Table 4.1](#), given that day's exchange rate of 1.25 CAD/USD.

By expressing exchange gains in terms of multiple currencies, we have not altered the *substance* of what has been recorded. We have only changed the *method* of recording.

Note that each line of [Table 4.2](#) is *balanced*, independently of the current exchange rate. Indeed, the entries of each currency are balanced separately: we have CAD 60 = CAD 180 – CAD 120, and also USD 100 = USD 100.

4.2. Accounting with currency trading accounts

An account that is denominated as a difference of multiple currencies, like the "exchange gain" account in [Table 4.2](#), is known as a *currency trading account*. The following is an exact copy of [Table 4.2](#), except that we have renamed the "exchange gain" account as "USD trading":

Table 4.3: Foreign currency accounting with a currency trading account

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	USD Trading (trading)
Jan 1	Balance (1 USD = 1.20 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 2	Balance (1 USD = 1.30 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 3	Balance (1 USD = 1.25 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120
Jan 4	Balance (1 USD = 1.15 CAD)	CAD 60	USD 100	CAD 180	USD 100 – CAD 120

It is important to note that a currency trading account does not represent *assets*. The purpose of a currency trading account is to calculate foreign exchange gains and losses. By convention, gains are recorded positively and losses negatively, and therefore a currency trading account is a kind of *income* account.

Currency trading accounts are different from any other account in that they are denominated in two (or more) currencies. Typically, one of these currencies is your "reference" currency, and the other one is in the usual terminology called a "security". Thus, in the above example, we may say that the USD trading account holds "securities" of USD 100, which are offset by CAD 120 in negative "currency".

A currency trading account reflects currency exchanges made in the past, and in an ideal situation where exchange rates never change, its balance is supposed to be zero. However, as the exchange rate fluctuates, so

does the value of the currency trading account, and such fluctuations correspond exactly to currency gains and losses. At the end of an accounting period, one may move the actual ("realized") gain or loss to an income or expense account denominated in the reference currency.

Let us now illustrate the use of a currency trading account by redoing the example from [Sections 2](#) and [3](#). We replace the "currency gain" account by a "currency trading" account between U.S. dollars and Canadian dollars.

Table 4.4: Foreign currency accounting with a currency trading account

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Food (expense)	USD Trading (trading)
Jan 1	Opening balance	CAD 200	USD 0	CAD 200	CAD 0	USD 0 – CAD 0
Jan 2	Exchange (1 USD = 1.20 CAD)	– CAD 120	+ USD 100			+ USD 100 – CAD 120
	Balance	CAD 80	USD 100	CAD 200	CAD 0	USD 100 – CAD 120
Jan 3	Buy food (1 USD = 1.30 CAD)		– USD 40		+ CAD 52	– USD 40 + CAD 52
	Balance	CAD 80	USD 60	CAD 200	CAD 52	USD 60 – CAD 68
Jan 5	Exchange (1 USD = 1.25 CAD)	+ CAD 75	– USD 60			– USD 60 + CAD 75
	Balance	CAD 155	USD 0	CAD 200	CAD 52	USD 0 + CAD 7
Jan 7	Buy food	– CAD 20			+ CAD 20	
	Balance	CAD 135	USD 0	CAD 200	CAD 72	USD 0 + CAD 7

Note that transactions that do not involve currency conversions (like "buy food" on January 7) are entered in the usual way for double entry accounting. Transactions that do involve currency conversions (like "exchange" on January 2 and 5 or "buy food" on January 3) must go through the trading account. Conceptually, when exchanging Canadian dollars for U.S. dollars, we are moving Canadian dollars "into" the trading account, and an equivalent amount of U.S. dollars "out of" the trading account. However, there is a sign reversal, so money moved "into" the trading account is given a negative sign and money moved "out of" it is given a positive sign. [FIXME: is there a standard convention for signs on trading accounts?]

Comparing the use of a currency trading account (as in [Table 4.4](#)) to the other accounting methods shown in [Tables 2.3](#) and [3.2](#), we note that the need for any periodic adjustments (as in [Table 2.3](#), January 3 and 5) or for periodic re-valuation of assets (as in [Table 3.2](#), January 3 and 5) has been eliminated. Indeed, the value of the trading account is "self-adjusting", because it is sensitive to exchange rates. Since no manual adjustments or re-valuations are needed, the question of *when* to record them becomes a moot point.

Another consequence of the self-adjusting nature of currency trading accounts is that there is no penalty for entering or changing transactions retroactively. Unlike in the SSAP 20 system, where retroactive transactions have an effect on future transactions and balances, in the currency trading account system, retroactive currency transactions only affect future balances (and not future transactions). In this respect, their effect is no different from other retroactive transactions.

Also note that, when using the currency trading account method as in [Table 4.4](#), there is no need to keep a separate "tally" or "inventory" of foreign currency assets, as we had to do using the SSAP 20 method in [Table 3.2](#) (green column). Instead, the foreign currency assets are held directly in an account.

Another advantage of the currency trading account method is that it provides some flexibility in tracking income and expenses from currency exchange. For many applications, it will be sufficient to have just a single currency trading account. But it is also possible to have several such accounts, for instance, one

account for each foreign currency, or one account for each branch of your business, or one account for each customer. This makes it possible to keep track of foreign exchange gains and losses attributed to different sources.

There is no reason that a single currency trading account has to be limited to two currencies. It is possible to set up a currency trading account to account for gains and losses from trades between three or more currencies. The balance in such an account is expressed as a sum or difference of amounts from each of the participating currencies.

Here is a summary of the main advantages of using currency trading accounts for multi-currency accounting:

- Foreign currency transactions are entered in accounts that are denominated directly in the foreign currency. Thus, the account sheet contains *source* information.
- Accounts are balanced; all transactions and balances respect the accounting equation.
- Foreign exchange gains and losses do not need to be calculated on a continuous basis; instead, they can be calculated whenever convenient (e.g. once per accounting period).
- Transactions can easily be entered or changed retroactively, without causing any change in future transactions.
- Foreign exchange gains and losses from different sources can be tracked independently (see [Section 4.4](#) for an example).
- There is no "reference" currency and thus no inherent asymmetry in this accounting method.

The most important conceptual facts to remember about currency trading accounts are the following:

- The purpose of a currency trading accounts is not to perform conversions, but to calculate gains and losses.
- The value of a currency trading account does not change during transactions, but in the time intervals between transactions.

4.3. Equivalence with SSAP 20

The accounting method using currency trading accounts is mathematically equivalent to methods such as SSAP 20. By this, I mean that it can be used to generate financial reports that are identical to those generated under SSAP 20. Indeed, suppose we were to generate a financial report from the data in [Table 4.4](#) at the end of January 3. We have recorded the following balances at the end of January 3:

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Food (expense)	USD Trading (trading)
Jan 3	Balance	CAD 80	USD 60	CAD 200	CAD 52	USD 60 – CAD 68

Using the exchange rate of 1.30 CAD/USD in effect on January 3, we translate all balances to CAD. Note that $USD\ 60 = CAD\ 78$, and therefore $USD\ 60 - CAD\ 68 = CAD\ 10$. Here is the "translated" balance sheet:

Date	Description	Canadian Cash (asset)	U.S. Cash (asset)	Initial Capital (capital)	Food (expense)	USD Trading (trading)
Jan 3	Balance	CAD 80	CAD 78	CAD 200	CAD 52	CAD 10

Keeping in mind that a currency trading account is a special kind of income account representing a foreign exchange gain, the translated balance sheet is *exactly identical* to the end-of-day balances for January 3 obtained from the SSAP 20 method that are shown in [Table 3.2](#).

Therefore, after translation to the "reference" currency, both accounting methods yield identical data and generate identical financial reports.

4.4. Example: tracking exchange gains and losses by source

Suppose you are a Canadian company selling goods to U.S. customers. Your books are mostly kept in Canadian dollars, but you invoice your customers in USD. Your customers are given 10 days to pay after you send them an invoice, and some additional time may pass until you receive and deposit their USD check in your CAD bank account. Foreign exchange gains and losses will occur if the exchange rate changes between the invoice date and the date the payment is received and converted to CAD.

Suppose that, for some reason, you would like to track foreign exchange gains and losses on a per-customer basis (for example, you want to determine whether it would be worthwhile to tighten some customer's payment terms). This can easily be done by using one currency trading account for each customer.

Table 4.5: One currency trading account per customer

Date	Description	Cash (asset)	U.S. Accounts Receivable (asset)	Income (income)	Trading Customer 2 (trading)	Trading Customer 1 (trading)
Jan 1	Opening balance	CAD 0	USD 0	CAD 0	USD 0 – CAD 0	USD 0 – CAD 0
Jan 2	Invoice Customer 1 (1 USD = 1.20 CAD)		+ USD 100	+ CAD 120		+ USD 100 – CAD 120
	Balance	CAD 0	USD 100	CAD 120	USD 0 – CAD 0	USD 100 – CAD 120
Jan 5	Invoice Customer 2 (1 USD = 1.30 CAD)		+ USD 200	+ CAD 260	+ USD 200 – CAD 260	
	Balance	CAD 0	USD 300	CAD 380	USD 200 – CAD 260	USD 100 – CAD 120
Jan 7	Payment Customer 1 (1 USD = 1.25 CAD)	+ CAD 125	– USD 100			– USD 100 + CAD 125
	Balance	CAD 125	USD 200	CAD 380	USD 200 – CAD 260	USD 0 + CAD 5
Jan 20	Payment Customer 2 (1 USD = 1.15 CAD)	+ CAD 230	– USD 200		– USD 200 + CAD 230	
	Balance	CAD 355	USD 0	CAD 380	USD 0 – CAD 30	USD 0 + CAD 5

The trading account balances on January 20 show that your interaction with customer 1 resulted in a foreign exchange gain of CAD 5, whereas your interaction with customer 2 resulted in a foreign exchange loss of CAD 30. The fact that a particular foreign currency interaction is "finished", i.e., there are no outstanding currency exchanges, is reflected in the fact that the USD balance in the corresponding currency trading account is 0.

The same method can be used to track foreign exchange gains and losses associated with any division of sources that you care to keep track of. For example, you might want to separate exchange gains and losses from different parts of your business, or from different countries (this works even when several countries share the same currency).

5. Realized and unrealized gains and losses

So far, I have discussed a method for calculating gains and losses from currency transactions (and by extension, from transactions with other types of variable-priced assets, such as barrels of oil or shares of stock). Such gains and losses are also called *capital gains and losses*. In practice, it is often important to distinguish between *realized* and *unrealized* gains and losses. One area where this distinction is relevant is taxation: usually, only realized gains are taxable. The exact determination of whether a particular gain or loss is realized or unrealized is often tricky, and the rules vary from country to country. In this section, I will briefly discuss these issues, and show how they can fit into the multi-currency accounting method.

In general terms, when you own a variable-priced asset, market fluctuations will cause the value of the asset to change over time. These changes of value are called *unrealized* gains and losses as long as you still own the item. When you sell the item, they turn into a *realized* gain or loss. While the basic principle is simple, the precise rules for when a gain/loss is realized depend on details of the tax code and accounting practices. In particular, it is not always obvious how to determine the realized gain or loss when you sell only part of an asset.

The calculation of *unrealized* gains and losses can sometimes require guesswork, because it depends on the fair market value of an asset that has not yet been sold. For assets that are unique and not traded on a daily basis, such as houses, it is often impossible to determine their precise value, and it is usually better to account for such items at their original purchasing price, rather than trying to re-evaluate them all the time. Some other types of assets, such as currencies and stock, are easier to evaluate. On the other hand, *realized* gains and losses can always be calculated exactly, because when an item is sold, its exact sales price is known.

5.1. An example from real estate

Let us start by considering a simple example from real estate. Suppose there are two houses on your street, called House A and House B. The houses are very similar, but they are on opposite sides of the street. Let us assume that the two houses have equal or similar market values at any given time.

Suppose on January 1, you have \$500,000. On February 1, you buy House A for \$100,000. On March 1, you buy House B for \$120,000. On April 1, your real estate agent estimates that, in the current market, each house is worth \$115,000. At this point, the market price of your two houses is \$10,000 more than what you paid for them, so you have an unrealized capital gain of \$10,000. More specifically, your unrealized gain from House A is \$15,000, and your unrealized loss from House B is \$5,000. On May 1, you sell House A for \$105,000, and you estimate that House B, which you keep, is also worth \$105,000. At this point, you have a *realized* gain of \$5,000 from House A (and you may have to pay taxes on this amount, minus any expenses, depending on the applicable tax laws). At the same time, you also have a total *unrealized* loss of \$15,000 from House B. On June 1, you decide to sell House B for \$110,000. You now no longer have any unrealized gains or losses. There is a realized loss of \$10,000 from your sale of House B. Together with the realized gain of \$5,000 from House A, this results in an overall realized loss of \$5,000 for the year.

Here is how the above scenario can be accounted for in the trading account model. The units are "CAD" for Canadian dollars, and "HOU" for houses.

Table 5.1: Realized and unrealized gains and losses: first real estate example

Date	Description	Canadian Cash (asset)	House A (asset)	House B (asset)	Initial Capital (capital)	House Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 500,000	HOU 0	HOU 0	CAD 500,000	HOU 0 – CAD 0	CAD 0
Feb 1	Buy House A	– CAD 100,000	+ HOU 1			+ HOU 1 – CAD 100,000	
	Balance	CAD 400,000	HOU 1	HOU 0	CAD 500,000	HOU 1 – CAD 100,000	CAD 0
Mar 1	Buy House B	– CAD 120,000		+ HOU 1		+ HOU 1 – CAD 120,000	
	Balance	CAD 280,000	HOU 1	HOU 1	CAD 500,000	HOU 2 – CAD 220,000	CAD 0
Apr 1	Balance	CAD 280,000	HOU 1	HOU 1	CAD 500,000	HOU 2 – CAD 220,000	CAD 0
May 1	Sell House A	+ CAD 105,000	– HOU 1			– HOU 1 + CAD 105,000	
	Realized gain					– CAD 5,000	+ CAD 5,000
	Balance	CAD 385,000	HOU 0	HOU 1	CAD 500,000	HOU 1	CAD 5,000

						– CAD 120,000	
Jun 1	Sell House B	+ CAD 110,000		– HOU 1		– HOU 1 + CAD 110,000	
	Realized loss					+ CAD 10,000	– CAD 10,000
	Balance	CAD 495,000	HOU 0	HOU 0	CAD 500,000	HOU 0 – CAD 0	– CAD 5,000

A number of observations are in order. First, unrealized gains and losses are implicitly recorded in the trading account, and thus fluctuate automatically as that account is evaluated at different times. For example, on March 1, the two houses are valued at \$240,000, and therefore the balance of the trading account (converted to CAD) is \$20,000, corresponding to the unrealized gain up to that day. On April 1, the two houses are valued at \$230,000, and therefore the trading account balance is worth only \$10,000.

Second, when an unrealized gain or loss becomes realized (on May 1 and June 1), this is witnessed by an explicit transaction, namely by moving the corresponding amount from a trading account to an income account. The important point is that realized gains and losses are not calculated implicitly by the accounting system, but must be explicitly stated by a bookkeeper or accountant. I will argue that this is a desirable and necessary feature. As the below examples will show, the calculation of realized gains in many cases depends on local laws and accounting practices, and it is therefore reasonable that it should be resolved by a human being. Of course the calculation can also be automated, but this should then be done by a customizable rule.

There is the question of *when* a realized gain or loss should be recorded. In the above example, and in the examples following below, I recorded it immediately on the date of the sale. However, if the number of transactions is large, it might be more convenient to calculate and record realized gains or losses less frequently. For example, one might decide to do so only once at the end of each accounting period.

If we had been very pedantic, we would have used two different units of "currency" for houses, for example HOA for quantity of House A, and HOB for quantity of House B. This is because the two houses, while similar, are not entirely interchangeable (after all, they are on different sides of the street!) Had we done this, we would have also created two separate trading accounts for HOA and HOB, similar to what was done in [Table 4.5](#). In this case, the realized gain for House A would have been the *total* content of the trading account for House A on May 1 (which would have been denominated in pure CAD on that day). Similarly, the realized loss for House B would have been the total content of the trading account for House B on June 1. In this way, it is possible to avoid doing a separate "side calculation" for determining realized gains and losses, at the expense of keeping a larger number of separate trading accounts.

5.2. Importance of the "identity" of assets

To see why things can get more complicated, let us change the above real-estate scenario slightly. Suppose you had decided to sell House B on May 1 and House A on June 1, rather than the other way around (after all, the houses are very similar). Then on May 1, you would have sold House B for \$105,000, resulting in a realized *loss* off \$15,000 for House B (compared to a realized *gain* of \$5,000 in the scenario where you sold House A on this day). On the other hand, on June 1, you would have sold House A for \$110,000, resulting in a realized gain of \$10,000 for House A (compared to a realized loss of \$10,000 in the original scenario). Here are the accounts for the updated scenario. (The changes between [Tables 5.1](#) and [5.2](#) are shown in red).

Table 5.2: Realized and unrealized gains and losses: second real estate example

Date	Description	Canadian Cash (asset)	House A (asset)	House B (asset)	Initial Capital (capital)	House Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 500,000	HOU 0	HOU 0	CAD 500,000	HOU 0 – CAD 0	CAD 0
Feb 1	Buy House A	– CAD 100,000	+ HOU 1			+ HOU 1 – CAD 100,000	
	Balance	CAD 400,000	HOU 1	HOU 0	CAD 500,000	HOU 1 – CAD 100,000	CAD 0

Mar 1	Buy House B	- CAD 120,000		+ HOU 1		+ HOU 1 - CAD 120,000	
	Balance	CAD 280,000	HOU 1	HOU 1	CAD 500,000	HOU 2 - CAD 220,000	CAD 0
Apr 1	Balance	CAD 280,000	HOU 1	HOU 1	CAD 500,000	HOU 2 - CAD 220,000	CAD 0
May 1	Sell House B	+ CAD 105,000		- HOU 1		- HOU 1 + CAD 105,000	
	Realized loss					+ CAD 15,000	- CAD 15,000
	Balance	CAD 385,000	HOU 1	HOU 0	CAD 500,000	HOU 1 - CAD 100,000	- CAD 15,000
Jun 1	Sell House A	+ CAD 110,000	- HOU 1			- HOU 1 + CAD 110,000	
	Realized gain					- CAD 10,000	+ CAD 10,000
	Balance	CAD 495,000	HOU 0	HOU 0	CAD 500,000	HOU 0 - CAD 0	- CAD 5,000

At the end of the year, your total realized loss for the two houses is \$5,000, regardless of which of the two scenarios you follow. So assuming your tax year ends on December 31, it does not matter which house you sold first. However, things are very different if your tax year ends on May 15, or if you wait several years before selling the second house. In this case, the difference for the outgoing tax year is between a gain of \$5,000 (in the first scenario) and a loss of \$15,000 (in the second scenario). **Therefore it matters which house you sold first, even if the houses are identical!**

5.3. The FIFO, LIFO, and adjusted cost base methods

Fortunately, in real estate transactions, there is never any question about which of several houses you sold. No matter how similar House A and House B are, they still have different addresses. However, things are not as simple if you own several units of an asset, such as shares of stock, that are actually indistinguishable from each other. In accounting terminology, such indistinguishable assets are called "fungible" assets. Let us redo the above example, but using shares instead of houses. We will also shorten the numbers by three zeros.

Suppose on January 1, you have \$500. On February 1, you buy one share of some company for \$100. On March 1, you buy another share for \$120. On April 1, you check the value of your shares, and you find that they are worth \$115 each, so their total value is \$230, or \$10 more than you spent to buy them. You have an unrealized capital gain of \$10. On May 1, you decide to sell one share for \$105. What is your realized gain at this point?

The answer is: it depends! This is precisely where the rules differ in different countries. There are three common cases.

- Some countries (such as [Australia](#)) use the *first-in-first-out (FIFO)* rule. Under this rule, the share sold on May 1 will be considered to be the share that was bought the longest ago, i.e., the one that was bought on February 1. The sale therefore results in a realized gain of \$5. This corresponds to the first real-estate scenario above (where the first house bought is the first house sold).
- Some countries use the *last-in-first-out (LIFO)* rule. As the name suggests, under this rule, the share sold on May 1 will be considered to be the share that was most recently bought, i.e., the one that was bought on March 1. The sale therefore results in a realized loss of \$15. This corresponds to the second real-estate scenario above.
- Some countries (such as Canada) use a method called the *adjusted cost base method* or the *weighted average method*. Under this method, the realized gain or loss is computed based on the average cost of all the shares currently owned. Thus, one would consider the total number of shares you own (2), and the total amount you spent on them (\$220, also called the *cost base*). The average cost of the shares is therefore \$110. The sale of one share for \$105 on May 1 will therefore be considered to create a

realized loss of \$5. After this, the cost base of \$220 is adjusted by subtracting the sold shares (\$105) and also subtracting the realized losses (\$5). The result is \$110, which is called the new *adjusted cost base*. This represents the total historical cost of the remaining shares after the sale (not their current value!), and will be used as the cost base for the next sale.

From the above descriptions of the FIFO and LIFO methods, it is clear that their implementation in the multi-currency accounting model looks almost exactly like the two real estate examples in [Tables 5.1](#) and [5.2](#). One just has to replace houses by shares. Of course, there will no longer be separate asset accounts for "House A" and "House B", but instead a single asset account for this company's shares. The FIFO example is shown in [Table 5.3](#), and the LIFO example is shown in [Table 5.4](#). To complete the examples, I have once again assumed that the last remaining share is sold on June 1 for \$110.

Table 5.3: Realized and unrealized gains and losses: FIFO

Date	Description	Canadian Cash (asset)	Shares (asset)	Initial Capital (capital)	Share Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 500	SHA 0	CAD 500	SHA 0 – CAD 0	CAD 0
Feb 1	Buy one share	– CAD 100	+ SHA 1		+ SHA 1 – CAD 100	
	Balance	CAD 400	SHA 1	CAD 500	SHA 1 – CAD 100	CAD 0
Mar 1	Buy one share	– CAD 120	+ SHA 1		+ SHA 1 – CAD 120	
	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0
Apr 1	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0
May 1	Sell one share	+ CAD 105	– SHA 1		– SHA 1 + CAD 105	
	Realized gain (FIFO)				– CAD 5	+ CAD 5
	Balance	CAD 385	SHA 1	CAD 500	SHA 1 – CAD 120	CAD 5
Jun 1	Sell one share	+ CAD 110	– SHA 1		– SHA 1 + CAD 110	
	Realized loss (FIFO)				+ CAD 10	– CAD 10
	Balance	CAD 495	SHA 0	CAD 500	SHA 0 – CAD 0	– CAD 5

Table 5.4: Realized and unrealized gains and losses: LIFO

Date	Description	Canadian Cash (asset)	Shares (asset)	Initial Capital (capital)	Share Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 500	SHA 0	CAD 500	SHA 0 – CAD 0	CAD 0
Feb 1	Buy one share	– CAD 100	+ SHA 1		+ SHA 1 – CAD 100	
	Balance	CAD 400	SHA 1	CAD 500	SHA 1 – CAD 100	CAD 0
Mar 1	Buy one share	– CAD 120	+ SHA 1		+ SHA 1 – CAD 120	
	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0

Apr 1	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0
May 1	Sell one share	+ CAD 105	– SHA 1		– SHA 1 + CAD 105	
	Realized loss (LIFO)				+ CAD 15	– CAD 15
	Balance	CAD 385	SHA 1	CAD 500	SHA 1 – CAD 100	– CAD 15
Jun 1	Sell one share	+ CAD 110	– SHA 1		– SHA 1 + CAD 110	
	Realized gain (LIFO)				– CAD 10	+ CAD 10
	Balance	CAD 495	SHA 0	CAD 500	SHA 0 – CAD 0	– CAD 5

For both the FIFO and LIFO methods, there is a non-trivial overhead involved in calculating realized gains and losses. These calculations require looking back at the history of shares acquired, figuring out which ones have already been sold, and then figuring out which of the remaining shares was the "first in" or "last in", respectively. A transaction for the realized gains or losses then needs to be created for each sale. There does not seem to be a convenient way to reduce this overhead in the FIFO and LIFO methods.

It turns out that the adjusted cost base method does not have such overhead. If we account for the above example by the adjusted cost base method, we end up with the following:

Table 5.5: Realized and unrealized gains and losses: adjusted cost base

Date	Description	Canadian Cash (asset)	Shares (asset)	Initial Capital (capital)	Share Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 500	SHA 0	CAD 500	SHA 0 – CAD 0	CAD 0
Feb 1	Buy one share	– CAD 100	+ SHA 1		+ SHA 1 – CAD 100	
	Balance	CAD 400	SHA 1	CAD 500	SHA 1 – CAD 100	CAD 0
Mar 1	Buy one share	– CAD 120	+ SHA 1		+ SHA 1 – CAD 120	
	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0
Apr 1	Balance	CAD 280	SHA 2	CAD 500	SHA 2 – CAD 220	CAD 0
May 1	Sell one share	+ CAD 105	– SHA 1		– SHA 1 + CAD 105	
	Realized loss (ACB)				+ CAD 5	– CAD 5
	Balance	CAD 385	SHA 1	CAD 500	SHA 1 – CAD 110	– CAD 5
Jun 1	Sell one share	+ CAD 110	– SHA 1		– SHA 1 + CAD 110	
	No realized gain (ACB)				CAD 0	CAD 0
	Balance	CAD 495	SHA 0	CAD 500	SHA 0 – CAD 0	– CAD 5

Note that, just like for FIFO and LIFO, each realized gain or loss has an associated transaction. However, in the adjusted cost base system, it is possible to calculate the amount of the realized gain or loss without

looking at the entire history of purchases and sales. If you check the example carefully, you will discover that the adjusted cost base at any given time is exactly equal to the CAD portion of the trading account. This is not a coincidence, as the trading account precisely mirrors the historical cost of shares. The realized gain or loss of a sales transaction can therefore be calculated simply by looking at the trading account balance just prior to the sale. Of course, this only works if there is a separate trading account for each kind of shares owned.

Of course, I am not trying to convince you to use the adjusted cost base method of accounting. You have to use whichever method is prescribed by the laws in your country!

In some countries, realized capital gains are taxed at different rates depending on how long you have held the asset. In such countries, you should have more than one income account for realized gains, for example, one for realized gains taxable at the lower rate, and the other for realized gains taxable at the higher rate. It is clear that one cannot use the adjusted cost base method in such situations, because there would be no notion of "how long" you have held the asset that you are selling. Instead, you should probably be using the FIFO method.

To further illustrate the adjusted cost base method, it may be useful to consider a more complicated example. Suppose you start with \$1000 on January 1. On January 2, you buy 100 shares of some stock at \$8 per share. On January 3, you sell 50 shares at \$9 per share. On January 4, you buy 50 shares at \$10 per share. On January 5, you sell 80 shares at \$11 per share. Using the adjusted cost base method, the realized gains and losses are computed on a daily basis as follows:

Table 5.6: Realized and unrealized gains and losses: second ACB example

Date	Description	Canadian Cash (asset)	Shares (asset)	Initial Capital (capital)	Share Trading (trading)	Realized Gain (income)
Jan 1	Opening balance	CAD 1,000	SHA 0	CAD 1,000	SHA 0 – CAD 0	CAD 0
Jan 2	Buy 100 shares (\$8/share)	– CAD 800	+ SHA 100		+ SHA 100 – CAD 800	
	Balance	CAD 200	SHA 100	CAD 1,000	SHA 100 – CAD 800	CAD 0
Jan 3	Sell 50 shares (\$9/share)	+ CAD 450	– SHA 50		– SHA 50 + CAD 450	
	Realized gain (ACB)				– CAD 50	+ CAD 50
	Balance	CAD 650	SHA 50	CAD 1,000	SHA 50 – CAD 400	CAD 50
Jan 4	Buy 50 shares (\$10/share)	– CAD 500	+ SHA 50		+ SHA 50 – CAD 500	
	Balance	CAD 150	SHA 100	CAD 1,000	SHA 100 – CAD 900	CAD 50
Jan 5	Sell 80 shares (\$11/share)	+ CAD 880	– SHA 80		– SHA 80 + CAD 880	
	Realized gain (ACB)				– CAD 160	+ CAD 160
	Balance	CAD 1030	SHA 20	CAD 1,000	SHA 20 – CAD 180	CAD 210

Therefore, under the adjusted cost base system, at the end of January 5, you have had \$210 in realized gains for the year, and you will still own 20 shares at an adjusted cost base of \$180.

It should be remarked that the discussion of accounting rules such as FIFO, LIFO, and adjusted cost base not only applies to shares, but also to other types of assets, most notably inventory. When a company sells its oldest inventory first (or assumes to do so for accounting purposes), it is effectively using FIFO. When it sells

its newest inventory first, it is using LIFO. If overall prices are rising, then the LIFO method will result in less reportable income (and therefore lower taxes). However, according to [this article](#), the LIFO method is frowned upon outside of the U.S., because it rarely makes sense to hoard old inventory.

While I have discussed capital gains and losses in the context of variably-priced *assets*, a similar discussion also applies to variable-priced *liabilities*. For example, suppose that you need to pay an invoice in a foreign currency. If that currency decreases in value, then the amount you have to pay also decreases (in terms of your own currency), and therefore you experience an unrealized gain. This turns into a realized gain when you actually pay the invoice. In some countries, like [Australia](#), gains and losses from assets must be reported separately from gains and losses from liabilities (in Australia, these are called "forex realisation events" of type 2 and 4, respectively). In such cases, one should have a separate income account for each type of realized gains and losses.

In conclusion, we can say that the multi-currency accounting model can be used to track realized gains and losses under a variety of accounting rules. Because realized gains and losses are entered as transactions, any specific rule, such as FIFO, LIFO, or adjusted cost base, can be implemented. Each of the systems requires a slightly different calculation to be done.

6. Translation issues

6.1. Independence of viewpoint

In [Section 2.3](#), we had briefly mentioned that the concept of a foreign exchange gain or loss apparently depends on one's choice of reference currency. A Canadian person who has assets of CAD 150 and USD 100 will perceive a shift in exchange rates from 1.20 CAD/USD to 1.25 CAD/USD as a gain, whereas an American person in the same situation will perceive such a shift as a loss.

In international business, such differences of viewpoint are more than just "perceptions"; they have real and important financial implications (for example, gains are taxed, but losses are not).

As an example, suppose that a Canadian company has a U.S. subsidiary. Suppose that the U.S. subsidiary has assets of CAD 150 and USD 100. For simplicity, also assume that the subsidiary has no transactions throughout the year, that the exchange rate on January 1 was 1.20 CAD/USD, and that the exchange rate on December 31 was 1.25 CAD/USD. Since the subsidiary company is American, its books are naturally kept in U.S. dollars. Here is what the books would look like, using the accounting method with a currency trading account:

Table 6.1: Books of the U.S. subsidiary company

Date	Description	CAD Cash (asset)	USD Cash (asset)	Initial Capital (capital)	CAD Trading (trading)
Jan 1	Opening balance	CAD 150	USD 100	USD 225	CAD 150 – USD 125
Dec 31	Closing balance	CAD 150	USD 100	USD 225	CAD 150 – USD 125

When preparing its financial statement for the fiscal year, the U.S. subsidiary "translates" the opening and closing balances into USD, using the exchange rate of 1.20 CAD/USD for January 1, and the rate of 1.25 CAD/USD for December 31. Here is the translation:

Table 6.2: Books of the U.S. subsidiary company, translated to USD

Date	Description	CAD Cash (asset)	USD Cash (asset)	Initial Capital (capital)	CAD Trading (trading)
Jan 1	Opening balance	USD 125	USD 100	USD 225	USD 0
Dec 31	Closing balance	USD 120	USD 100	USD 225	– USD 5

Date	Description	CAD Cash (asset)	USD Cash (asset)	Initial Capital (capital)	CAD Trading (trading)
Jan 1	Opening balance	CAD 150	CAD 120	CAD 270	CAD 0
Dec 31	Closing balance	CAD 150	CAD 125	CAD 270	CAD 5

The financial report of the Canadian subsidiary will look like this:

Initial capital (Jan 1): CAD 270
Plus Gains: + CAD 5
Final capital (Dec 31): CAD 275

The parent company can incorporate this statement directly into its books, since it is already in Canadian dollars.

What this example illustrates is *independence of viewpoint*, an important accounting principle. The end result, from the point of view of the parent company, is the same, no matter whether the subsidiary keeps track of its equity in USD or in CAD. In the first case, the parent will report capital gains of CAD 11.25, which are partially offset by currency exchange losses of CAD 6.25. In the second case, the parent will report currency exchange gains of CAD 5.00. The net gain, in each case, is the same.

In summary, it does not matter whether one first translates the books to a new currency and then prepares a financial report, or whether one first prepares a report and then translates the report to a new currency, adding a capital gain or loss on the initial capital.

In practice, the nationality of a subsidiary can of course make a difference, for example in the area of taxation. In the above example, the U.S. subsidiary reports a net loss, and is therefore not taxed in the U.S. (however, its Canadian parent is taxed in Canada). The Canadian subsidiary reports a gain, and it is therefore potentially taxed in Canada (although in the case we considered here, it probably won't be taxed, as the gain is unrealized). Any such taxes are partially offset by lower taxes for the Canadian parent.

6.2. Reference currency vs. Equity currency

Note that [Tables 6.1](#) and [6.3](#) make use of trading accounts, and thus, they use an accounting method that treats multiple currencies in a symmetric way. There is no "reference currency" in the sense of SSAP 20, because not all accounts are kept in a single currency.

Nevertheless, there is of course a difference between [Table 6.1](#) and [6.3](#). An American company will typically record its equity (capital, income, and expenses) in USD, while a Canadian company will do so in CAD. Therefore, each company will favor a particular currency.

However, this particular asymmetry is caused by the way companies choose to denominate their accounts, and is not built into the accounting *system*. There is no reason why, in principle, some company could not keep some of its income, expense, or capital accounts in USD and others in CAD. The accounting method with currency trading accounts would have no problem accommodating this. However, the preparation of single-currency financial reports would be somewhat trickier in this case, as capital gains and losses would have to be added to the reports for foreign-denominated capital accounts (this can be done by the method shown in [Section 6.1](#)).

7. Specific software

This section provides some pointers to multiple currency accounting support in specific accounting software packages.

7.1. GnuCash

GnuCash has built-in support for multiple currencies, but prior to version 2.3.9, its double-entry accounting followed the method described in [Section 2.2](#) ("the wrong way"), with currency gains and losses calculated by the reporting system. Proper currency trading accounts could be set up manually, but were not supported by the exchange rate system, so each transaction had to be calculated by hand. As of version 2.3.9, support of currency trading accounts has been added, but is not enabled by default. See [Multiple currency accounting in GnuCash](#) for information on how to enable it, and also for information on the older "manual" method.