What rate of return to allow? Do we truly understand it?

Arnoud W.A. Boot and Jeroen Ligterink[*] / April 25, 2011

1. Introduction

Many utility companies and infrastructure companies that have a natural monopoly position are subject to regulation to prevent them from charging excessive tariffs. Examples include TenneT's electricity transmission system and KPN's fixed-line network, which other parties must be able to use through interconnection, and Schiphol Airport's aviation activities that airlines depend on.

Part of the regulation entails determining the maximum return on invested capital that can be passed on. This is referred to as rate-of-return regulation. The maximum rate of return on the invested capital is – in principle – determined such that it matches the rate of return required by the providers of capital (i.e. the investors in the capital market). The underlying idea is that the rate of return on an activity to which a dominant position appears to be attached, must be limited to a level which just matches the required return of investors in the capital market. It should be neither too high nor too low.

In this contribution we discuss how the normative rate of return should be determined and we assess the underlying methodology and theory. Both the determination of the normative rate of return and the application of the theory of corporate finance on which it is based are highly vulnerable to misconceptions. Our aim is to contribute to a better understanding and hence to improvements in the future practice of the rate-of-return regulation.

We also address the limitations associated with the rate-of-return regulation and its methodology. What exactly is the purpose of this approach and to what extent is it implemented correctly in practice? Although the subject of rate-of-return regulation may, at first glance, appear to be somewhat technical and perhaps even boring, the many fascinating aspects that it encompasses must not be overlooked. Of interest, for example, is the very basic question as to the precise purpose of regulating the rate of return on the capital that is employed. As already indicated, the rate-of-return regulation is designed to prevent the abuse of a dominant position. At the same time, however, it should provide the right incentives for inducing optimal investment behaviour. That is, if the normative rate of return is set too high, it may induce overinvestment; if it is set too low, it may induce underinvestment. More in general, the regulated party should have optimal incentives to operate efficiently and effectively. The normative rate of return that is imposed and other elements of the 'regulatory package' should not be at odds with this.[1]

The investment behaviour is an interesting issue because it directly affects the capital that is employed and therefore the total compensation (in €) on the invested capital; this compensation being the normative rate of return times the total (regulated) capital employed (i.e. invested capital). An expanding asset base which does not benefit the end users can therefore be just as damaging for them as an excessively high rate of return.

In section 2, we discuss the role of rate-of-return regulation within the cost-plus and price-cap approaches that are common in the regulation of natural monopolies. It is important for the reader to understand that defining a norm for the rate of return plays a similar role in both approaches. In section 3, we deal with the basic principles of rate-of-return regulation building on the core insights from the field of corporate finance. A familiar concept that will be introduced here is Beta as a measure of risk and therefore as a key determinant for the normative rate of return. The higher the Beta, the higher the rate of return required by investors and therefore the rate of return that will be allowed. Section 4 shows that the normative rate of return can be considered as a weighted average cost of capital (WACC) of debt and equity financing.

In section 5, we discuss the limitations of corporate finance theory. In particular, we consider the extent to which the exclusive focus on Beta as a measure of risk is justified. Section 6 discusses implementation issues relating to the WACC and Beta. Section 7 focuses on the distinction between single-till and dual-till regulation. This distinction is of importance in situations in which there are related activities that may or may not fall within the scope of regulation. Section 8 concludes with a few sobering remarks about the stunning importance that the rate of return or WACC methodology has acquired in practice.

2. Two approaches to the regulation of tariffs

The two standard approaches to the regulation of tariffs are the cost-plus approach and the price-cap approach. In the price-cap approach utility companies may increase prices each year in accordance with, for example, the change in the consumer price index corrected for expected annual changes in efficiency and/or quality. In the cost-plus approach, the price is set such that it covers the company's allowed costs plus a maximized rate of return on its asset base.

Comprehensive and in-depth treatises can be written about the possible differences between the price-cap approach on the one hand, and the cost-oriented approach on the other. These differences are less relevant to this contribution, however, since we focus on the determination of the rate of return on invested capital, which is relevant to both approaches anyway. That is, in the cost-plus approach the return is explicitly added, while in the price-cap approach a remuneration is included for the invested capital that is tied up ('employed') in the activity. Nevertheless, we will indicate differences where they are important. If the price-cap approach is strictly applied, for example, usually more risk remains with the regulated company because it is more difficult to pass on (unexpected) changes in costs. The cost-plus approach is often more accommodating. This affects the normative rate of return on the invested capital, since more risk must be reflected in a higher allowable rate of return. In practice though, one must be extremely cautious when comparing the approaches because, when it comes to implementation, a great variety of hybrid elements are often included that blur the strict theoretical distinction.

3. Basic principles of rate-of-return regulation

How can one establish the appropriate allowable rate of return on invested capital? To answer this question, a general method is available which results from theories developed in the field of corporate finance. The method is derived from the Capital Asset Pricing Model (CAPM) which has been developed in the 1960s and is the

'workhorse' of corporate finance. Although the model itself requires specialised knowledge to understand, it is easy to explain in terms of use. The model indicates that the rate of return that investors require on an investment in a particular activity equals the (appropriate) risk-free rate plus a compensation for risk, a risk premium. This premium only includes non-diversifiable risk (measured by Beta). Diversifiable risk plays no role. Why this is the case, at least according to the underlying theory, is explained below.

Risk can be mitigated through diversification, for instance by simultaneously investing in different companies (or activities). Assuming, as is done in the field of finance, that investors are risk averse on average, diversification is a good thing because it reduces risk. In finance jargon, diversifiable risk is referred to as company-specific risk. This concerns company-specific events that do not play a role within a diversified portfolio of investments. That is, spreading investments across a multitude of companies makes this risk 'disappear'; setbacks for one company are compensated by windfalls for another. The important insight is that since investors can avoid being harmed by company-specific risk through diversification, they do not require compensation for it in terms of a higher rate of return. It does therefore not show up in the risk premium.

The non-diversifiable risk, also referred to as market risk, does require compensation. The degree of market risk associated with an investment can be measured by means of the familiar Beta. Market risk refers to the degree of sensitivity to the (world) economy, or the market as a whole. This risk cannot be avoided through diversification.[2] It must be borne by somebody. While everyone can individually opt to be free of market risk (e.g. by not investing, or by investing in activities that are immune to the aggregate state of the economy), we cannot altogether avoid it. Hence there will always be investors who must bear this risk. This should not come as a surprise, since together we must bear the fluctuations of the economy as a whole.

What does this mean for an individual investor? An individual investor can only avoid market risk by passing it on to someone else. This is the reason that compensation is required for bearing such risk, since other individuals will only be prepared to take it over (and therefore bear it), if there is a compensation for doing so. The more an asset moves with the market, the greater the market risk associated with it, which is reflected in the investment's Beta. A higher Beta means that an investment entails greater market risk, and thus that investors will require a higher risk premium. Beta is based on the idea that, because market risk is non-diversifiable, it has a price. In other words, because this risk does not simply disappear, someone must bear it, and that party will require compensation for doing so. In accordance with finance theory, this compensation takes the form of a risk premium in the rate of return required by investors. The risk premium is proportional to the investment's Beta, which is a measure of the quantity of market risk that the investment is subjected to.

Company-specific risk, on the other hand, 'melts like snow in the sun' in a diversified portfolio. Indeed, an investor can eliminate this risk through diversification. In keeping with the earlier wording, it is possible to avoid company-specific risk without passing it on to someone else. Hence, no compensation is needed; it can costlessly be avoided by everybody. Thus, the required rate of return is not affected by diversifiable risk.

We have now almost completed the full circle. The normative rate of return with respect to a regulated activity is determined by the degree of market risk associated with the activity concerned. The market risk is the reason for the risk premium that is added to the risk-free rate in the calculation of the required rate of return according to the CAPM. An activity free of market risk is therefore allowed to pass on a rate of return equal to the (well chosen) risk-free interest rate. For an activity with market risk, an additional risk premium applies which is proportional to the activity's Beta.

A complication is that sometimes a company simultaneously carries out regulated activities as well as other, non-regulated ones. A first step in understanding this complication is how to look at individual activities when they are jointly undertaken by one firm. An important result from the field of corporate finance is that the required rate of return for an activity must be defined solely on the basis of that particular activity's risk profile. By extension, it follows that if a regulated activity is carried out by a large conglomerate, the risk of the conglomerate as a whole should not affect the determination of the required rate of return on the regulated activity. The reason is that the risk of the conglomerate is affected by all the other activities that the conglomerate undertakes. A separate estimate must therefore be made for the normative rate of return on the regulated activity. That is, the question is how the investors view the activity concerned on its own merits. For finance wizkids, the core result in corporate finance pertaining to capital budgeting decisions applies: in determining the net present value of an investment project, the expected future cash flows of the project must be discounted against the required rate of return that is determined *solely* by the project's risk profile and not by the risk profile of the company as a whole.[3]

Things get more complicated if there are synergies between the company's regulated activities and other activities taken up by the same company. This makes it more difficult to identify the cash flows that can be attributed to the regulated activities. One such example is Schiphol Airport, where non-aviation activities, such as retail revenues, fall outside the scope of the regulation. For now, it is sufficient to note that it is precisely this kind of example that can lead to conflicts. Certain non-aviation activities at Schiphol Airport depend to a significant extent on the regulated aviation activities. For example, if there were no flights, there would be fewer retail customers. A question then is whether these retail revenues should not count – at least in part – as returns on the regulated activities (and thus indirectly lower the permitted return on the regulated activities). The other activities could also affect the actual all-in risk profile of the regulated activities. In addition, it might be difficult to attribute investments to one or another activity. These interrelationships are further discussed in section 7.[4]

4. Weighted average cost of capital (WACC)

The rate of return allowed under regulation is, in principle, determined by the rate of return that investors or providers of capital would require for the financing of such an activity. In the implementation of rate-of-return regulation, the normative rate of return is linked to a notional weighted average cost of capital (WACC). In layman's terms, it is assumed that the regulated activity is financed as a separate company by means of debt and equity. Assuming a certain (normative) debt to equity ratio, the holders of debt and equity require specific rates of return on their claims. So if we

know the debt to equity ratio and their respective required rates of return, we could compute the weighted average cost of capital or WACC. This is explained in the box below.

The essence of the WACC method

The weighted average cost of capital (WACC) is the weighted average of the required rate of return on debt (after tax) and the required rate of return on equity, with the market value of debt and the market value of equity as fractions of the total value as weights. The formula is as follows:

$WACC = K_d \times (1-T) \times (D/(D+E)) + K_e \times E/(D+E)$

in which:
K_d: cost of debt (required rate of return for debt holders);
T: marginal tax rate;
D: market value of debt;
E: market value of equity;
K_e: cost of equity (required rate of return for equity holders).

The WACC is usually determined on the basis of the Capital Asset Pricing Model (CAPM). The CAPM states that the required rate of return on an asset or security (debt or equity in this case) depends on the risk-free interest rate and a premium for the non-diversifiable risk, based on Beta as the measure of applicable market risk for that asset. The formula is as follows:

$R_i = r_f + ß_i (R_m - r_f)$

in which:
R_i: required rate of return on asset i, where R_i is either K_d or K_e;
r_f: risk-free interest rate;
$ß_i$: Beta of asset i, with i being debt or equity;
$(R_m - r_f)$: market risk premium.

The following variables are important for the determination of the WACC on the basis of the CAPM:
- Risk-free interest rate (r_f)
- Beta and the cost of equity (K_e= required rate of return on equity)
- Cost of debt (K_d)
- Leverage (gearing g; equals Debt / Regulatory Asset Base)

The determination of these parameters is far from trivial. A comprehensive appendix detailing the way in which the different parameters can be determined is available from the authors (in Dutch, see *www.accf.nl/**boot**/index.php?id=1*; the appendix also contains a bibliography of relevant works).

We have already stated that the required rate of return for the investors (i.e. providers of capital) depends on the risk profile of the activity to be regulated. The division of invested capital into its source, equity and/or debt, seems to be a somewhat unnecessary complication. Why not assume that the entire activity is (notionally) financed by equity only? In principle, doing so should be acceptable because, as Nobel Laureates Modigliani and Miller have taught us, the way in which something is financed is of secondary (if any) importance. That is the risk of an activity must be borne by the providers of capital, irrespective of the way in which the capital is

divided. Opting for two kinds of capital – equity and debt, for example – does not alter the total risk of the activity. The only thing that happens is that the risk is distributed in a particular way, allocating proportionally more risk to equity and proportionally less to debt, as debt has priority over equity.

The choice of capital structure is therefore primarily an issue of distribution, and is not of substantive importance. After all, the weighted average risk of equity and debt is equal to the risk of the activities on the left hand side of the balance sheet. And note that these activities are the source of risk. Risk originates on the left hand side of the balance sheet and gets distributed on the right hand side of the balance sheet. If we recall that providers of debt and equity both require a compensation proportional to the risk they bear, the distribution of risk should not matter for the total compensation that the firm has to pay, and therefore capital structure would not matter. The weighted average cost of capital (the WACC) should therefore – in principle – not be affected by changes in the ratio between equity and debt.

This is the key insight of Modigliani and Miller. There is considerable misunderstanding about this concept in practice, since it is thought that debt is 'cheaper' than equity. This is an incorrect way of thinking. Debt only appears to be cheaper because it carries less risk due to the priority it is given. This priority is at the expense of equity holders and makes debt look cheap and equity look expensive. Nevertheless, in regulatory practice, it has become conventional wisdom that a notional capital structure must indeed be chosen in which a certain (maximum) level of debt is chosen such that at least a minimum credit rating is guaranteed.[5] This limits the amount of debt. It is assumed in the calculations that this partial debt financing leads to a lower allowable rate of return.

In principle, this runs counter to the view of Modigliani and Miller. Modigliani and Miller stress that all the fuss about the ratio between equity and debt is merely an issue of distribution and not substance. In other words, it distributes the risk that is created on the left hand side of the balance sheet over these two forms of capital, but the total risk remains the same, and hence the total risk premium that needs to be paid to debt and equity holders together. The only real justification for the assumption that changing the ratio of equity to debt has an effect, is the unequal tax treatment of debt versus equity. Due to the tax deductibility of interest payments, debt is considered to be cheaper than equity. It is indeed the case that if a third party (the government) 'contributes' towards one kind of capital (debt) and not towards the other kind (equity), then financing by means of (more) debt results in a saving for the company. Consequently, there is a justification for the WACC approach.[6]

5. Is Beta taken too seriously?

While great for finance theorists, including ourselves, to see theoretical concepts like WACC and Beta gain such prominence in reality, some qualifications should be made. It can be questioned whether people are truly aware of the kind of theoretical world that must be created to develop Beta as *the* measure of risk. Among other things, it requires everyone to have the same (homogeneous) expectations about the future and specifically the asset or activity considered. Investors should also experience no 'frictions' in their investment decisions, i.e. transaction costs cannot exist. These criteria are obviously not met in the real world.

The use of Beta as the measure of risk requires that investors face no transaction costs because they must be able to spread their investments as optimally as possible without incurring costs and thereby remove the company specific risk through diversification. While in the real world transaction costs are a fact of life, one might say that investors could – in principle – via index funds and ETFs obtain diversification at relatively low cost.[7]

Would this mean that a company is indifferent to diversifiable risk – that is, to company-specific risk? Of course not. Every entrepreneur knows that one is just as likely, or perhaps even more likely, to run aground as a result of company-specific risk than as a result of market risk. This means that companies most certainly take into account company specific risk, and it therefore cannot be simply ignored. This could mean that a regulated activity which through the WACC receives compensation only for Beta, i.e. market, risk, faces insufficient investment incentives because the company concerned experiences other risks that make it reluctant to invest. Such risks could include, for example, fear of a change in regulations imposed by the government for which insufficient compensation is anticipated.

Empirical research shows that also investors take company-specific risk into account. There is no exclusive focus on market risk. The conclusion is that risk is important, and goes beyond market risk. Some modesty is therefore required to ensure that the theoretically watertight distinction between non-diversifiable market risk and diversifiable company-specific risk in the CAPM world is not interpreted in such absolute terms. The practical experience of risk is a more complex phenomenon.

6. Implementation of WACC and Beta

Apart from the considerations in the previous section, the implementation of WACC and Beta-like concepts is by no means a straightforward affair. The measurability of the parameters of the WACC and the determination of the Beta itself can lead to almost unsolvable problems. For example, what is the Beta of a particular activity? It must accord with how its rate of return will move together with the rate of return on the market as a whole in the future. But how do we know this? After all, it refers to future movements. Often beta is estimated on the basis of historical data. The assumption then is that the type of activity concerned will move with the market in much the same way as it did in the past. If the structure of the economy does not fundamentally change over time, for Beta estimation purposes historical data could give a meaningful direction for the future.

Another important parameter for the WACC is the market risk premium. The market risk premium stands for the compensation *per unit* of market risk that investors – in expected value sense – wish to receive on top of the risk-free rate as compensation for bearing such risk. The market risk premium is the remuneration for the risk that applies to a diversified investment in the market portfolio. The market portfolio is a 'basket' of all investments available in the market, weighted according to size. Typically, a particular broad equity market index is chosen for this purpose. The market portfolio has – by definition – a Beta of one.

What is the remuneration for the risk associated with the market portfolio? Once we know that, we can establish the required compensation (on top of the risk free rate) for a specific activity; that compensation equals the Beta of that activity times the market

risk premium. Again, we require a forward-looking estimate. The past could serve as a guideline for the future. For example, based on the past, we could conclude that, on average, the equity market performed 4 or 5 percentage point better each year than risk-free bonds.[8] From this, we could conclude that the market risk premium is approximately 4% or 5%. However, this line of reasoning confuses the past with the future. Why would this average continue to apply in the future?

The issue is, in fact, a more complex one than the relevance of the past for establishing the Beta. The reason is that two diametrically opposed views could be posed as to how the past could be interpreted for the future. Note that both approaches realize that a forward-looking estimate is needed. According to one view, if we have data for the past 60 years and the equity market performed, on average, x percentage point better each year than risk-free bonds, then this provides a meaningful estimate for the future. Apparently the dynamics in the market enforce such premium, and why would this suddenly be different in the future? After all, we do not have any other 'hard' information, so how could we justify any other estimate?

The problem is that the alternative view leads to diametrically opposite insights. It draws the conclusion that a high estimate of the risk premium based on the past must entail a low estimate for the future. So the future does not follow the past, but goes counter to it. The underlying line of reasoning is as follows.

A high estimate of the market risk premium based on the past means that high rates of return were achieved in the past. A reason for this could be that prices in the stock market have risen due to a decreasing risk aversion on the part of investors. For example, increasing (average) wealth over time might have enhanced the willingness of people to assume risks in their investments. Note that stock prices should have increased in value at the point in time when that decrease in risk aversion manifested itself. This relationship is not difficult to see: people are willing to pay more for the same shares when they are less risk averse.

The past therefore exhibits a high rate of return on the stock market (stock prices increased) but expected future rates of return are lower because of lower risk aversion (and higher starting share prices). This completes the circle. Due to a decrease in risk aversion, share prices have shot up in the past, and this is what can explain the high rates of returns on equity observed in the historical data. From this higher level, however, expectations with respect to future rates of return are correspondingly lower. After all, stock prices have reached a high level from which it is more difficult to achieve a high rate of return. Since investors have become less risk averse, they are willing to accept such lower expected future rate of return, and therefore a lower market risk premium. In this view, the high returns in the past imply lower expected returns in the future!

Thus, there are two possible views on the relation between past and future rates of return, and they are diametrically opposed. This has enormous consequences for the estimation of the WACC. The market risk premium is far from a trivial parameter. It must be noted, perhaps superfluously, that the discussion is not about the validity of the theory but, rather, about the way in which to get to estimates of key parameters in the WACC.

7. Mixture of regulated and non-regulated activities: dual-till versus single-till regulation

The interrelationship between regulated and non-regulated activities can give rise to considerable complications. In section 3, we stated that only the risk profile of the activity to be regulated must be considered when determining the normative rate of return. How does this apply when the cash flows are interrelated? For example, in regulating aviation activities of airports where investments in aviation can also benefit non-aviation activities, one is confronted with the question how to deal with these synergetic (interrelated) cash flows. This is the issue in dual-till versus single-till regulation. Under a single-till regime, all activities are regulated as part of one integrated bucket, whereas in a dual-till regime regulation applies exclusively to the 'monopoly' activities. In practice, hybrid variants frequently apply.

If a pure dual-till approach is adhered to, then a related activity remains outside the scope of regulation. It could be argued that this promotes entrepreneurship in the organisation because profits achieved outside the regulated part do not come under the rate-of-return regulation and therefore provide real benefits for the investors and, indirectly, for the managers (potentially higher pay) in the company. This can also be favourable to the regulated activity if it gains as a result, for instance, because such benefits provide an additional incentive to ensure that the regulated activity is performed according to high standards. Schiphol Airport, for example, could have a greater interest in exceptionally clean terminals (part of the regulated aviation activity) to ensure that passengers are more inclined to spend time at the airport in order to visit its shops (the non-regulated activity). The risk here is potential 'gold plating', i.e. that, at the expense of the regulated activity, more is invested in the regulated activity than is socially desirable (excessive cleaning or something else).[9]

A different concern is precisely the mirror image of the entrepreneurship referred to above. Although entrepreneurship sounds good, it can likewise mean that management allows itself to be too distracted by the non-regulated activity and starts ignoring the regulated activity. Countering such distractions could benefit the utility nature of the regulated activity. Another aspect is that entrepreneurship in the non-regulated activity can also mean opportunism of a kind that leads to high risks that impact (threaten?) the regulated activity. If this is indeed the case, the continued existence of a regulated activity can be put in jeopardy by the non-regulated activity. These concerns can be addressed by the single-till form of regulation, which serves to contain the degrees of freedom in both types of the activities simultaneously, possibly at the cost of less entrepreneurship.[10]

8. Conclusion

One way to interpret this contribution is that it is somewhat surprising how generally accepted the WACC approach is in the rate-of-return regulation. Many questions can be raised with respect to both the theoretical underpinnings and the practical implementation. The problem, however, is that there is no readily available alternative.

From a theoretical perspective, the WACC is important for inducing optimal investment incentives. But investment incentives are also heavily influenced by other events. A disaster such as the credit crisis that took place in 2007-2009 is a case in

point. It does not seem possible to sufficiently fine-tune a regulation method to deal with such extreme circumstances.[11]

Another legitimate concern about the WACC method is that regulated activities often require long-term investments, whereas remunerations based on the WACC method are usually fixed for shorter periods of time, often for one year or for a certain contract period with a user. The long-term nature of the investments together with the perceived risk from say changes in regulations (a form of sovereign risk) or other circumstances are perhaps not sufficiently taken into account.

One should also be aware that investment decisions are affected by opportunity costs in general. This concerns all scarce resources, not only the financial capital on which the WACC method focuses but also other resources like available land, the reach of management and so forth, and particularly the alternative uses of such resources. This latter aspect is of particular importance in situations in which, as discussed in section 7, there is a mixture of regulated and non-regulated activities in an organisation.

The overall conclusion must be that the popularity of rate-of-return regulation and the precision suggested by the WACC methodology are somewhat surprising. Considering the many pitfalls of the methodology, particularly in its implementation, great care should be taken and outcomes will be far from perfect, whatever the objectives.[12]

[1*] Arnoud Boot is Professor of Corporate Finance and Financial Markets at the University of Amsterdam. Jeroen Ligterink is Associate Professor in Finance at the University of Amsterdam.

[1] In most cases, many other instruments designed to contribute to a desired level of efficiency and effective operations exist simultaneously. Examples include instruments like yardstick competition and a broad range of requirements pertaining to (participation in determining) the investment policy.

[2] Unless we start trading with Mars and investors can allocate their investments to both the Earth and Mars. In this case, a degree of diversification of market movements on Earth would be possible by also taking up positions on Mars. After all, the market on Mars would certainly move somewhat differently than the one on Earth. We have not yet reached this stage, however.

[3] The underlying idea is that the finance doctrine stresses the importance of value additivity. Apart from synergies in cash flows, diversification within the company is not important. Investors in the market can accomplish this diversification themselves. Therefore, companies can simply specialise in what they are good at and must not themselves diversify. This is also connected to the exclusive focus on Beta risk. In section 5, we set out a number of qualifications regarding the exclusive role of the market risk within the finance doctrine.

[4] In the case of airports, this problem is dealt with in different ways by regulators. In certain countries, single-till regulation is applied, which essentially means that the activities that coincide with the aviation activity are regulated in a single bucket. In other countries, such as the Netherlands, dual-till regulation is applied, which means that aviation activities are in principle considered separately. In practice, a hybrid structure is almost always in place (see section 7).

[5] The underlying idea is that too much debt causes an excessively high probability of bankruptcy, and that this entails additional costs. In addition to the existence of taxes (see later), such bankruptcy costs are a second deviation from the world conceptualised by Modigliani and Miller.

[6] Also in the case of Schiphol Airport (100% government owned), the government tries to secure a degree of debt financing, since the prevailing tax legislation means that such financing creates a saving for the company. In the case of a government owned company, however, this is a bizarre game, since the amount that the company gains as a result accrues to the government as a shareholder, while, at the same time, the same government loses exactly the same amount in tax revenue.

[7] Note that investors often incur sizable transaction costs also with mutual funds. This is still one of the puzzles, how such expensive asset management industry can continue to exist.

[8] We do not expand on how to define the market portfolio. The market portfolio includes all risky assets. Often the (local?) equity market is chosen as representation for this broader market.

[9] The existence of non-regulated activities in a dual-till situation can also mean that a far lower tariff, and therefore a far lower actual return, is deliberately chosen for the regulated activity because, in one way or another, doing so benefits the non-regulated activity. This appears to be akin to offering very affordable hotel accommodation in Las Vegas in order to attract 'captured' customers to the hotel owned casino. In an average Las Vegas hotel, finding the exit has been made as difficult as possible and, insofar as the exit can even be found, it can only be reached after a host of obstructive gaming machines have been negotiated. Another point is that due care must be taken to avoid potential competition-related problems. These may arise if the non-regulated activity becomes inappropriately powerful relative to other providers of this activity that do not simultaneously benefit from having a monopoly on the regulated activity.

[10] Single-till regulation appears to be logical only for activities that are related to each other. However, risk can likewise originate from non-related activities and can lead to major financial or other problems that can damage the regulated activity. In this case, it is not illogical to set restrictions with respect to who is permitted to be the owner of a certain regulated activity.

[11] In the most extreme case, such a disaster could mean the bankruptcy of some companies with regulated activities. It is highly unlikely that the government will refrain from intervening in such a case. Consider, for example, Schiphol Airport. If that airport had faced permanently lower air traffic due to the financial crisis or permanent 'ash cloud' problems, what would have happened? The problems concerning the operation of the high-speed railway line between Amsterdam and Belgium are a different but nevertheless pertinent example in this regard. The operator of the initial lease (i.e. the winner in the bidding war for the right to exploit the new high speed rail connection) appears bankrupt. The government, however, might feel compelled to come to the rescue.

[12] The recent credit crisis is an interesting illustration of the practical and theoretical problems related to the WACC approach. The credit crisis gave rise to a (short term) increase in the market risk premium and the credit spread (i.e. the difference between the cost of debt and the risk free rate). What is the implication for the WACC to be used in regulation? Should it go up? We are of the opinion that one should be very careful to draw implications from these increases. Financial markets were frozen and companies adapted by reducing their investments, irrespective of the WACC.