

# Data Structures

## AVL Trees - 2

CS 225

Harsha Tirumala

October 3, 2025



UNIVERSITY OF  
**ILLINOIS**  
URBANA-CHAMPAIGN

Department of Computer Science

# Learning Objectives

Review why we need balanced trees

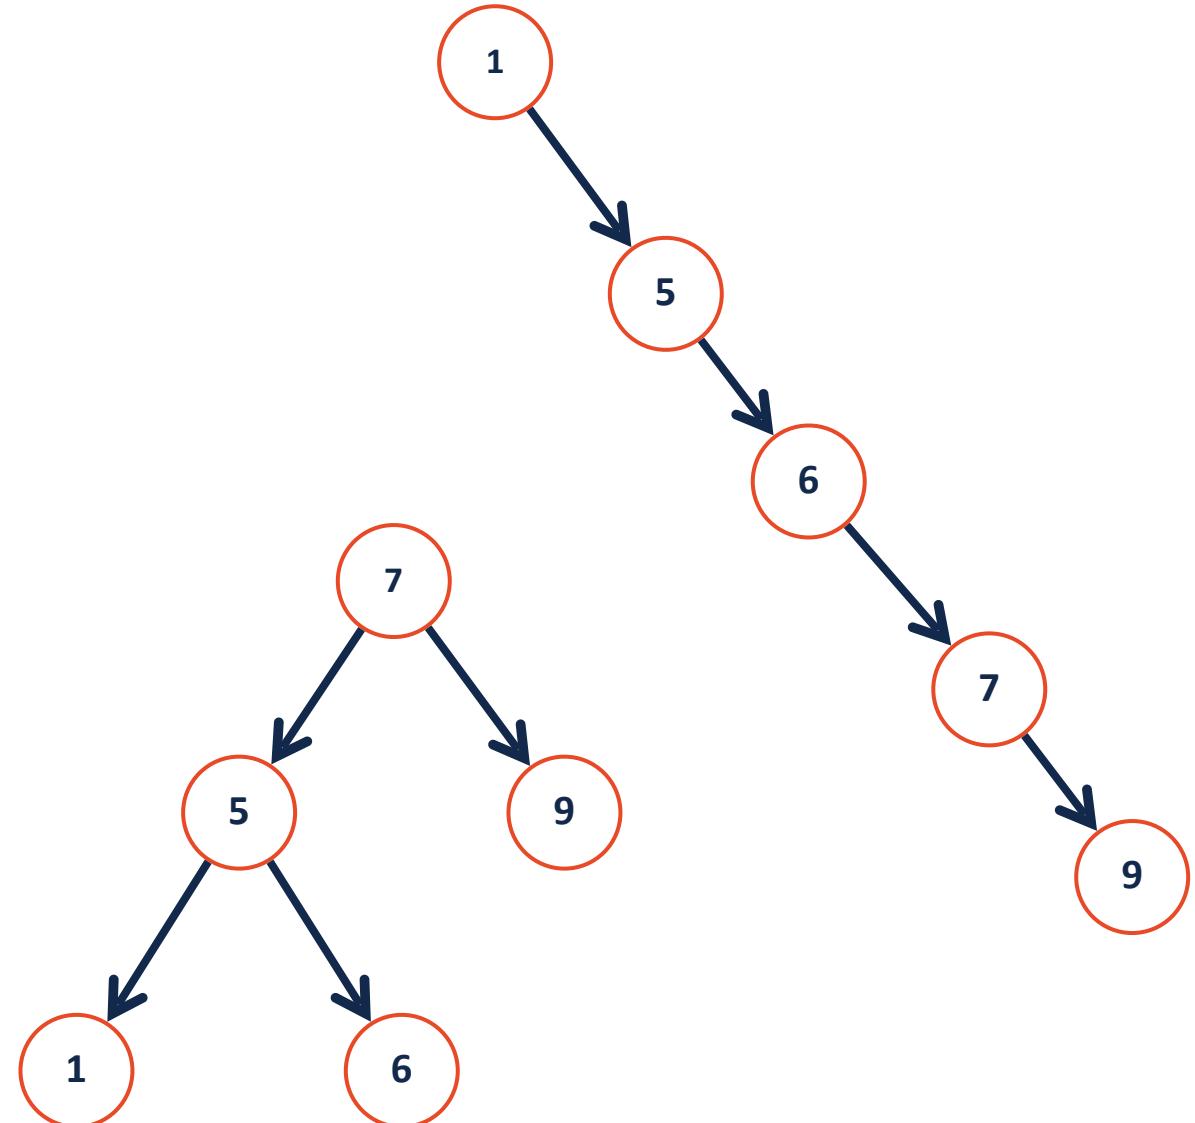
Review what an AVL rotation does

Review the four possible rotations for an AVL tree

Explore the implementation of AVL Tree

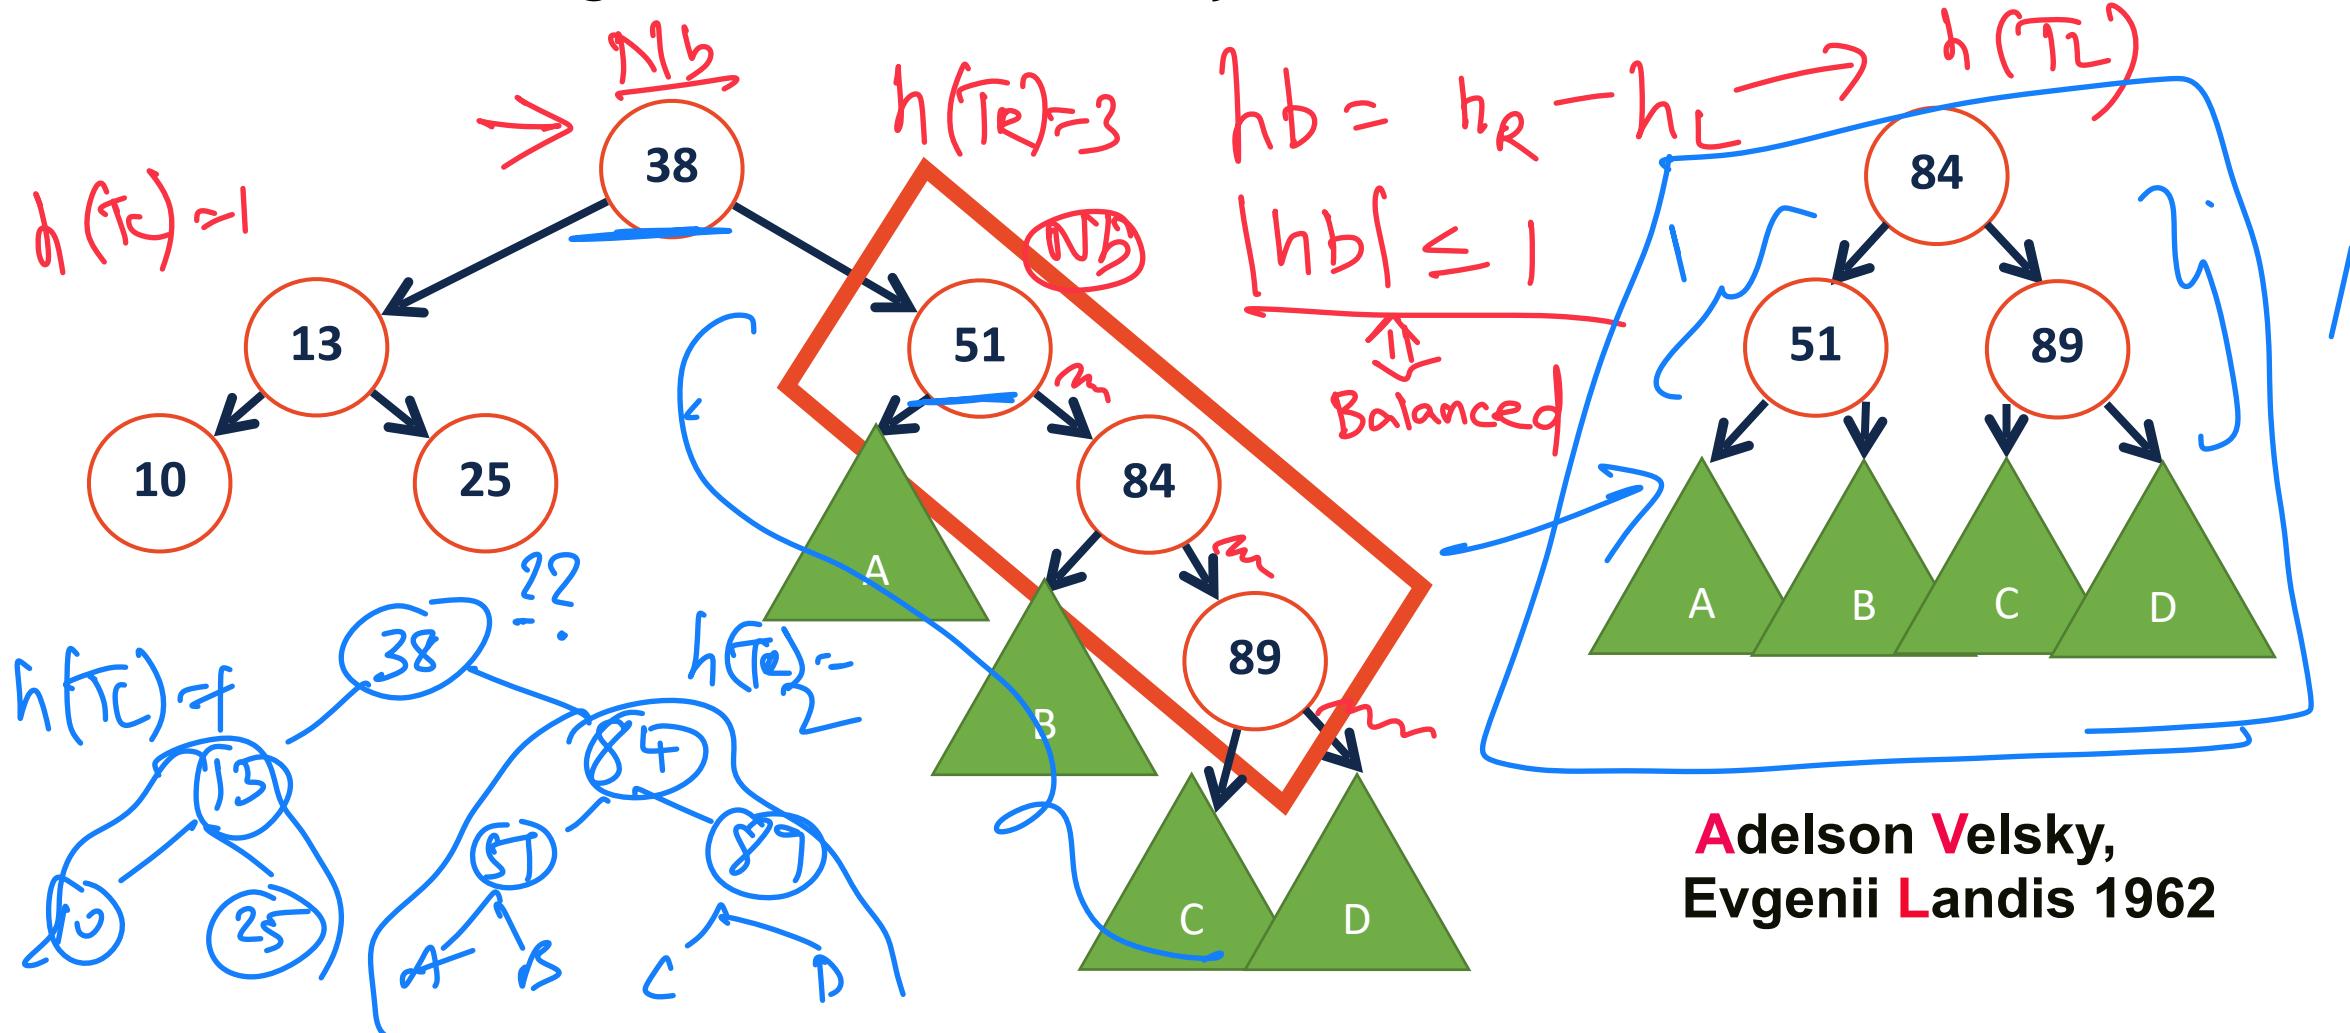
# BST Analysis – Running Time

	BST Worst Case
find	$O(h)$
insert	$O(h)$
delete	$O(h)$
traverse	$O(n)$



# AVL-Tree: A self-balancing binary search tree

Rather than fixing an insertion order, just correct the tree as needed!

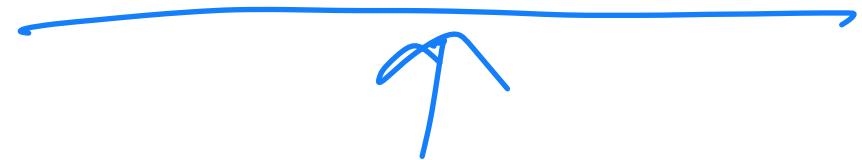


# BST Rotations (The AVL Tree)

We can adjust the BST structure by performing **rotations**.

These rotations, when used correctly:

1. Modify the arrangement of nodes while preserving BST property



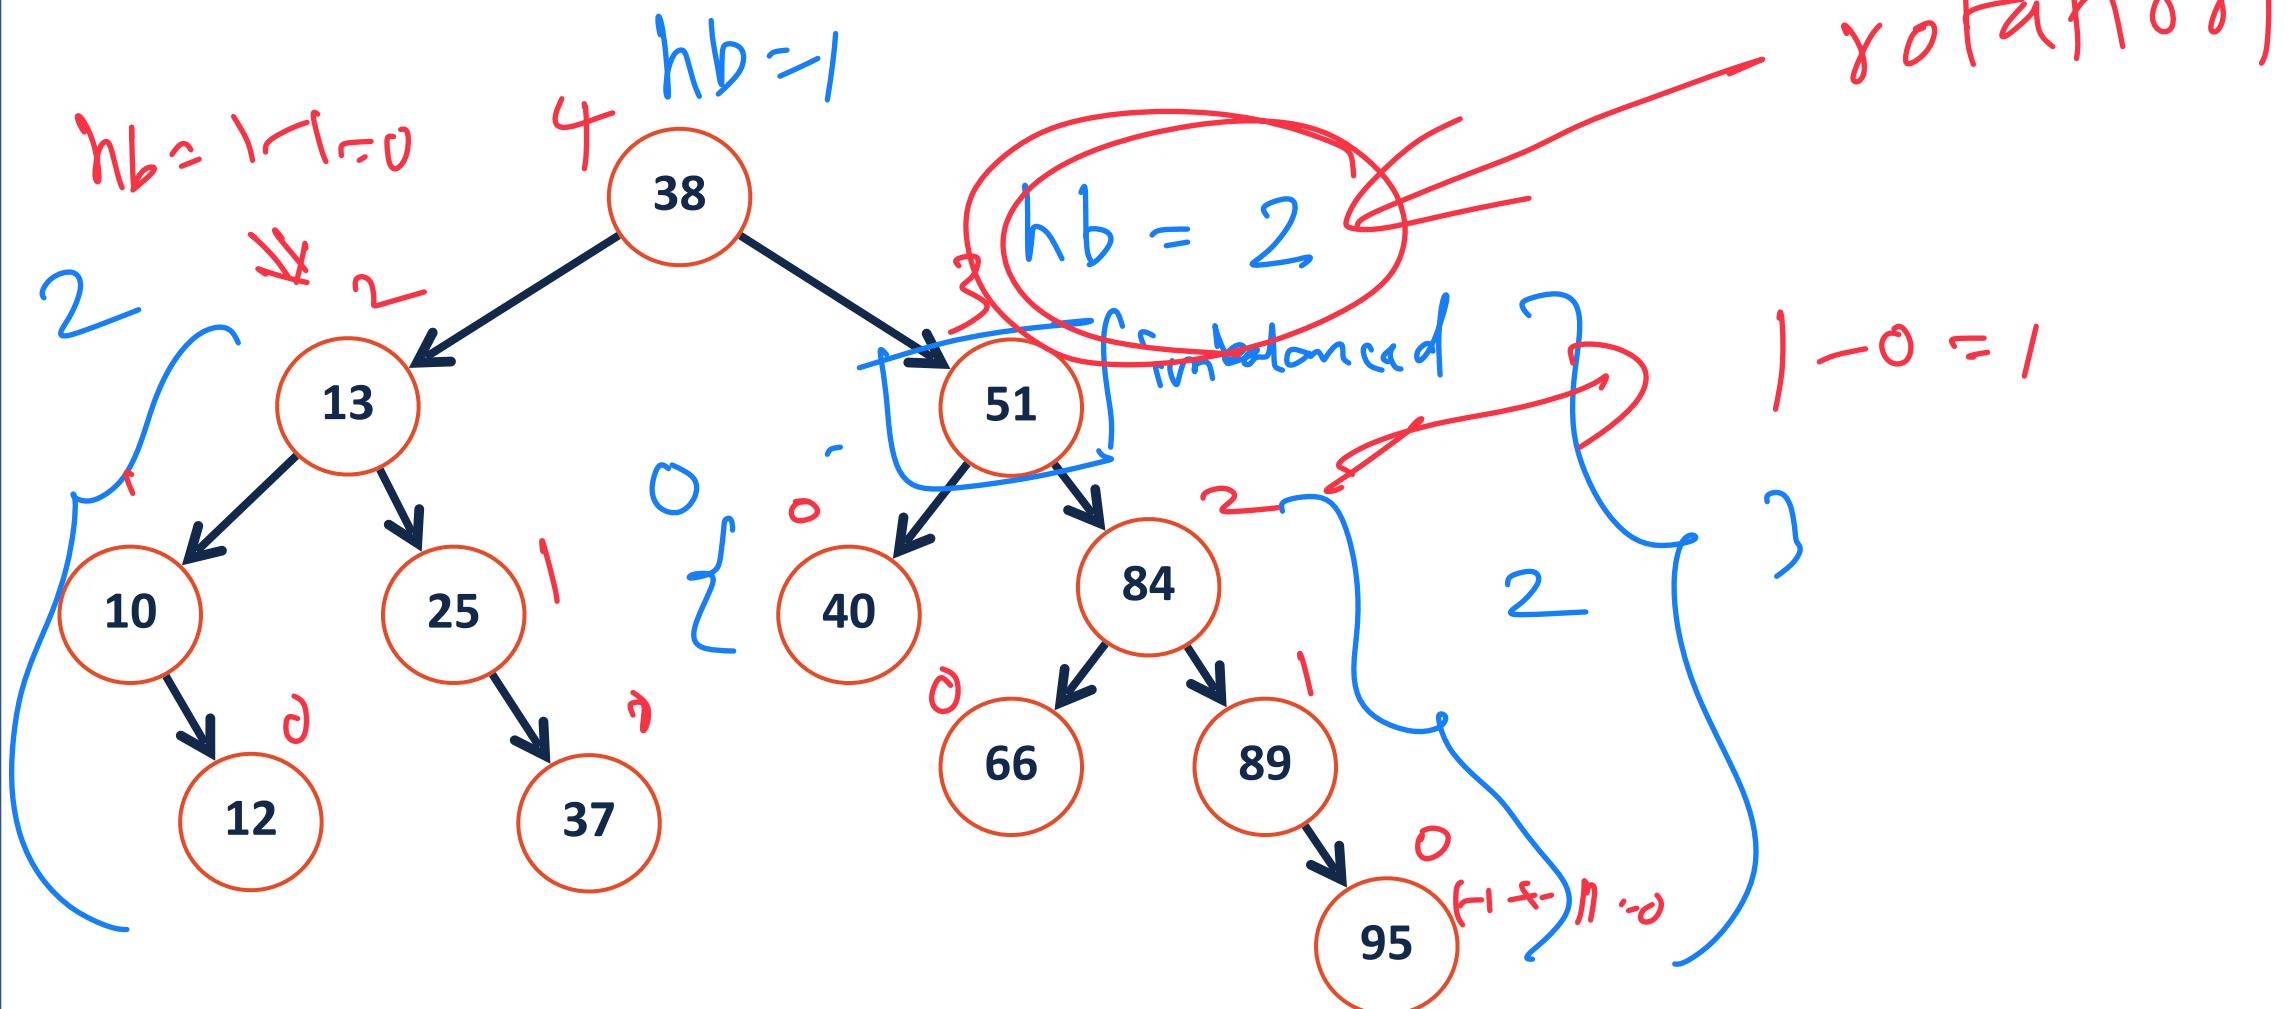
2. Reduce tree height by one

(Balance)

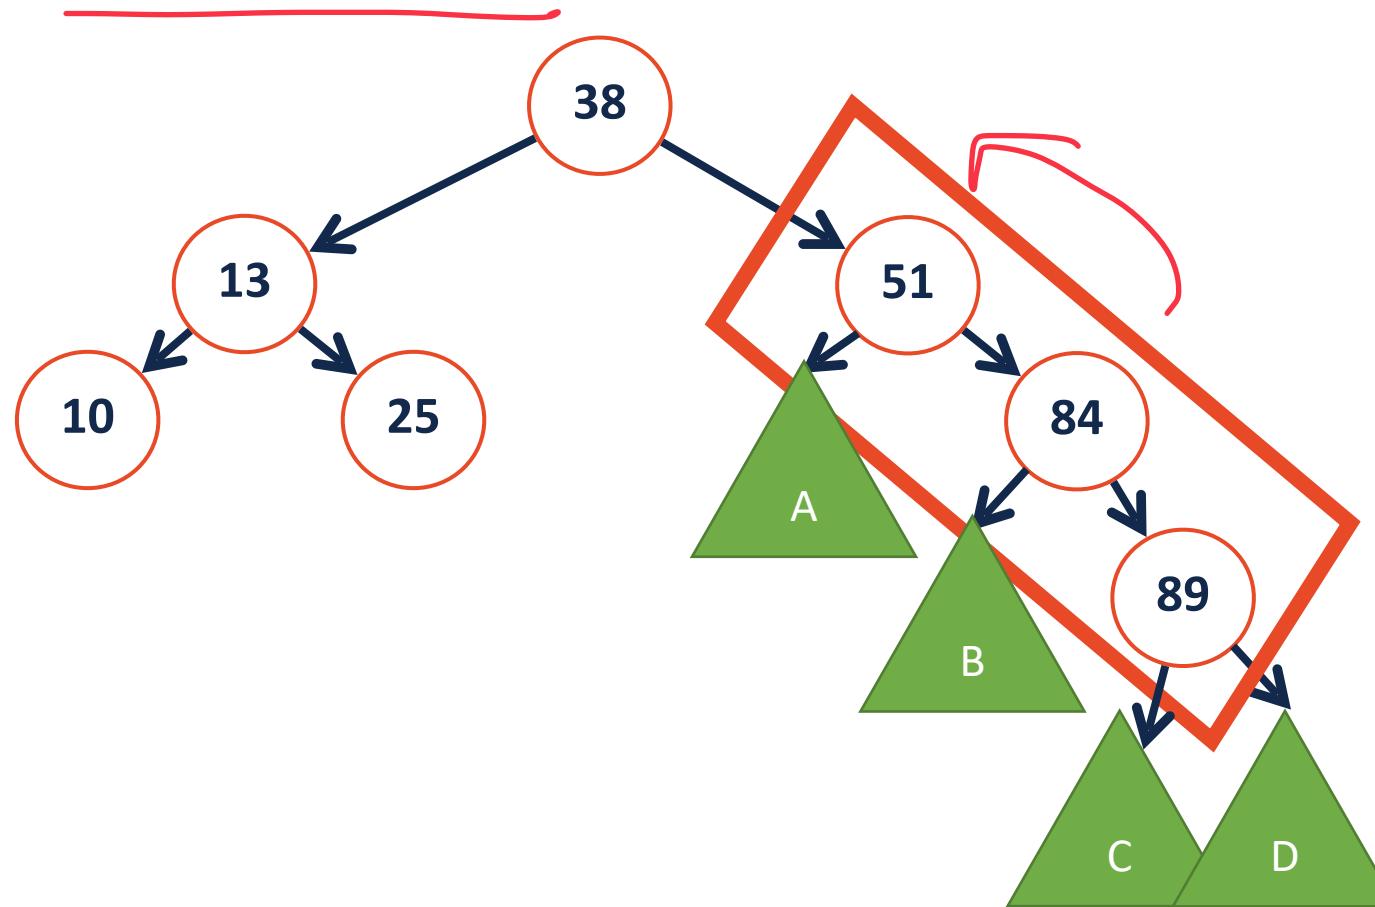


# BST Rotations (The AVL Tree)

To begin, lets find the imbalance in the following tree:

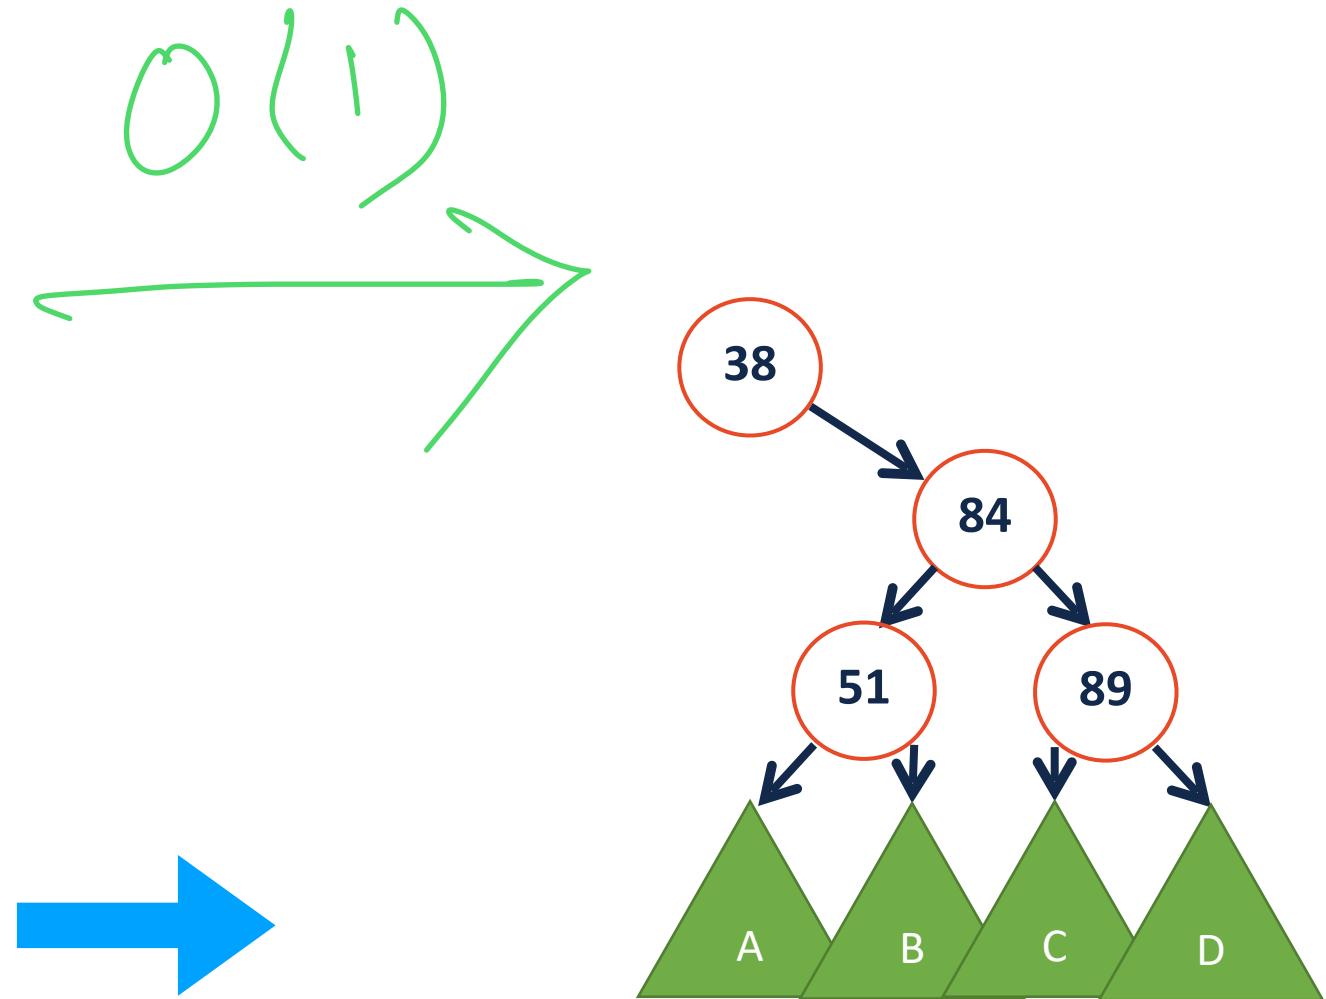
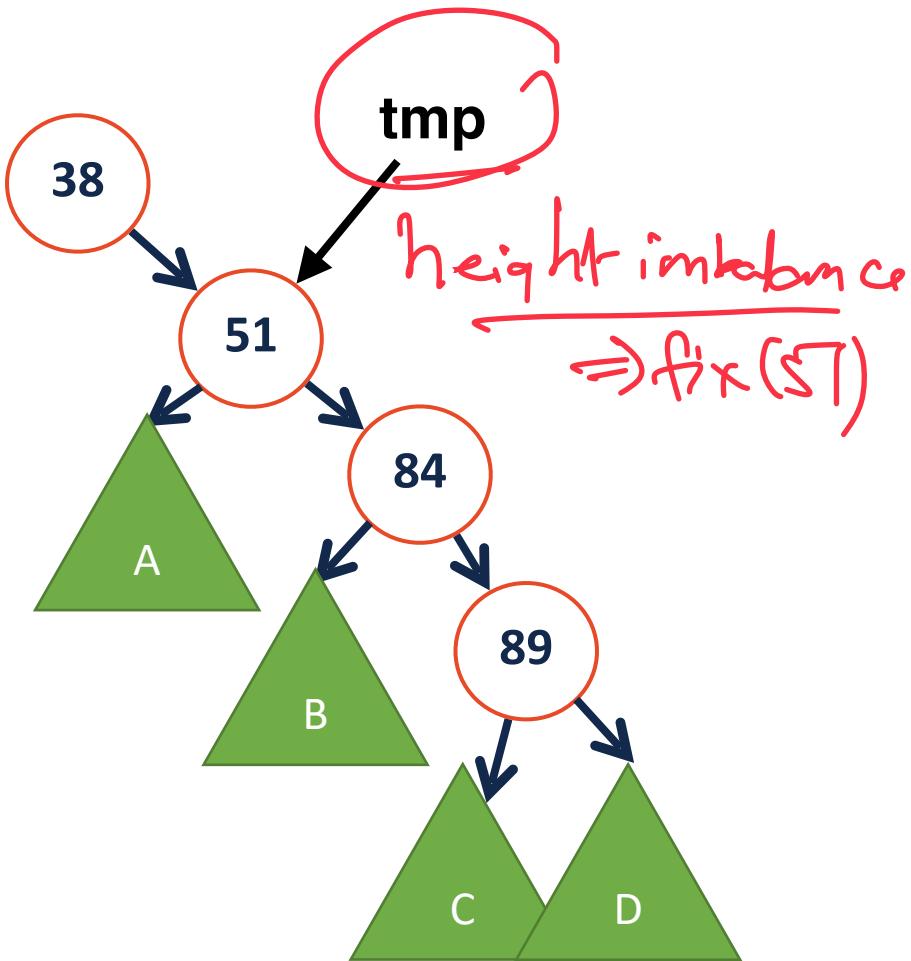


# Left Rotation



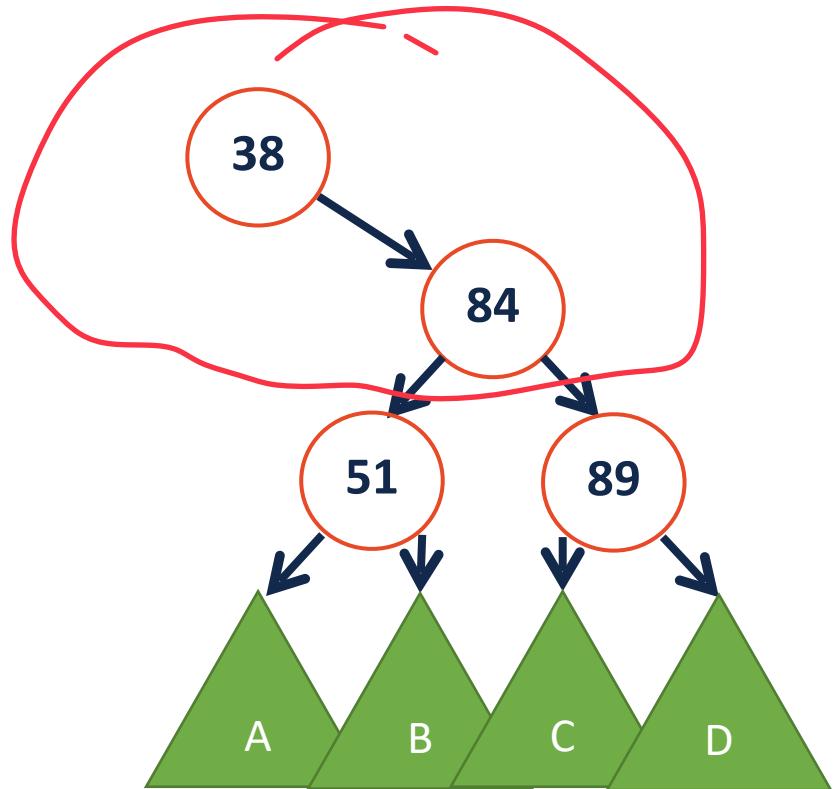
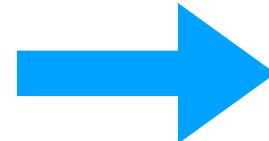
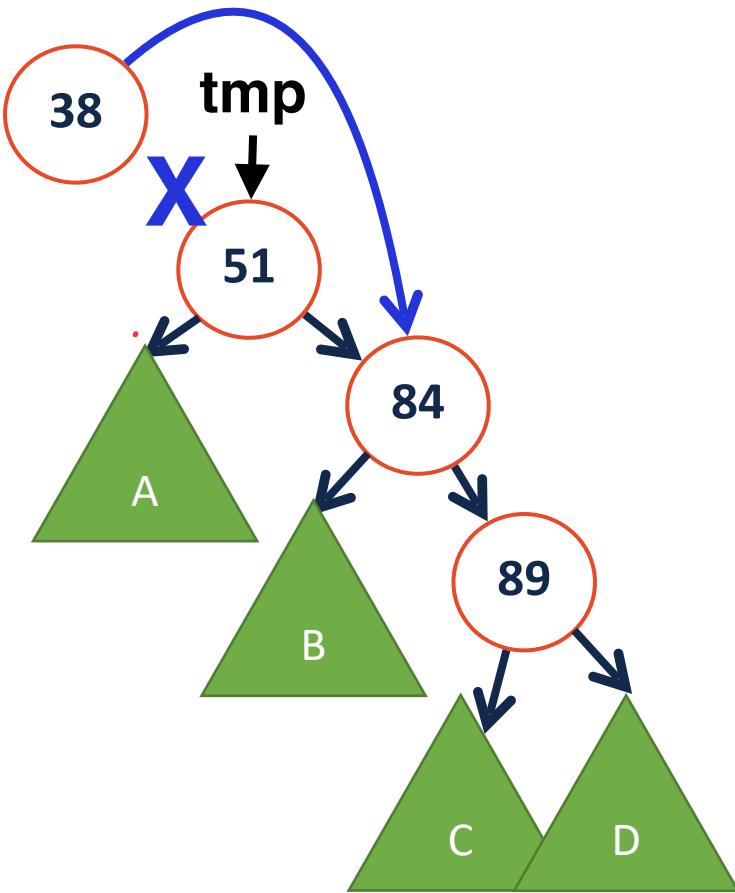
# Left Rotation

1) Create a tmp pointer to root



# Left Rotation

- 1) Create a tmp pointer to root
- 2) Update root to point to mid

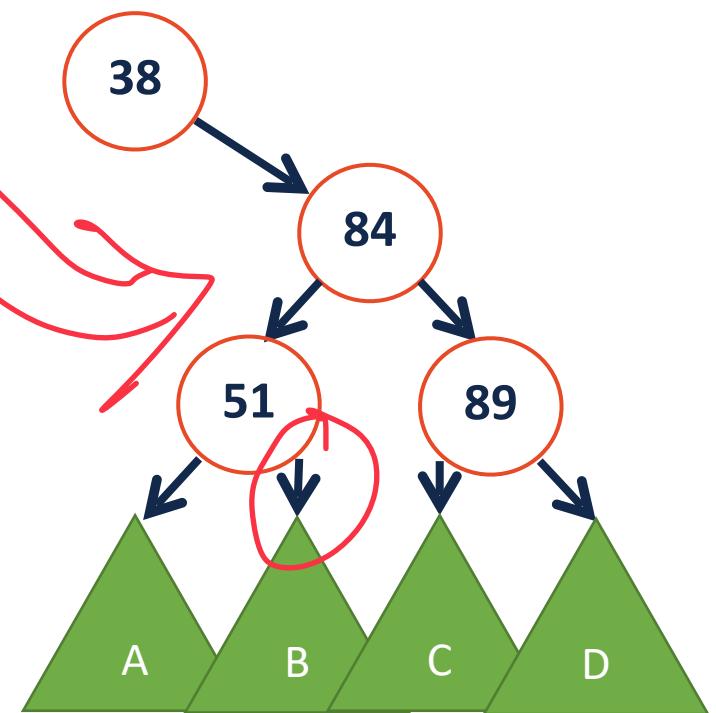
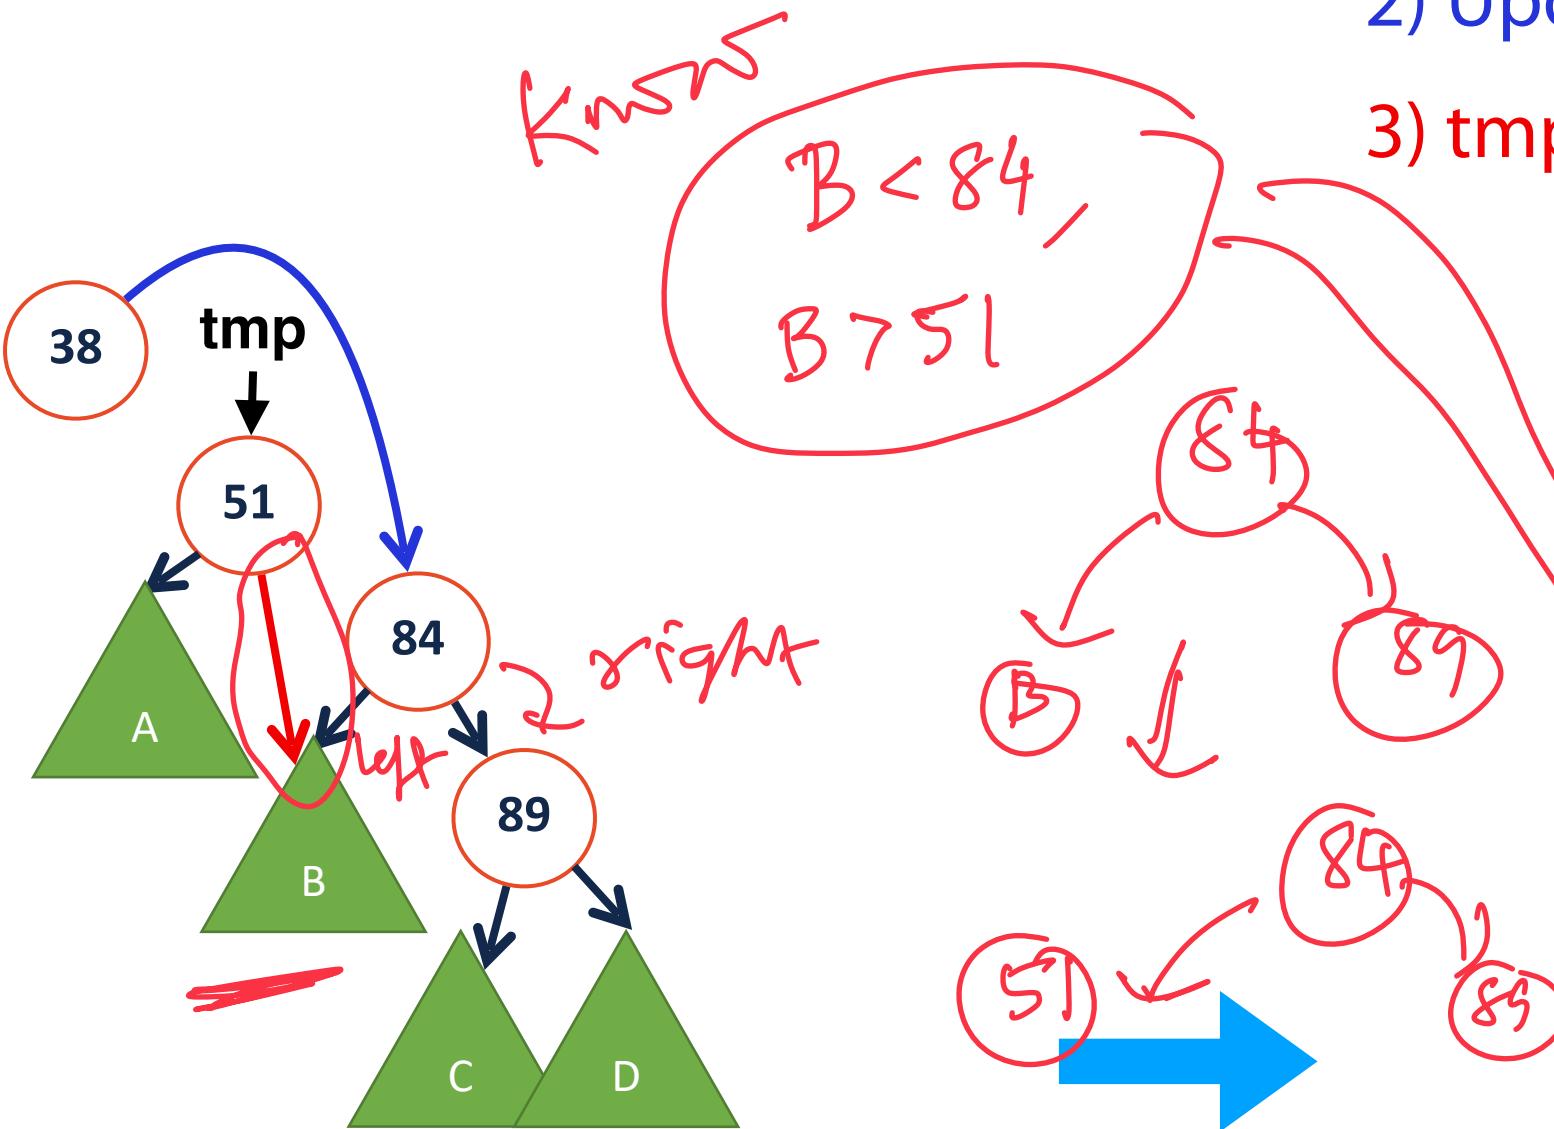


# Left Rotation

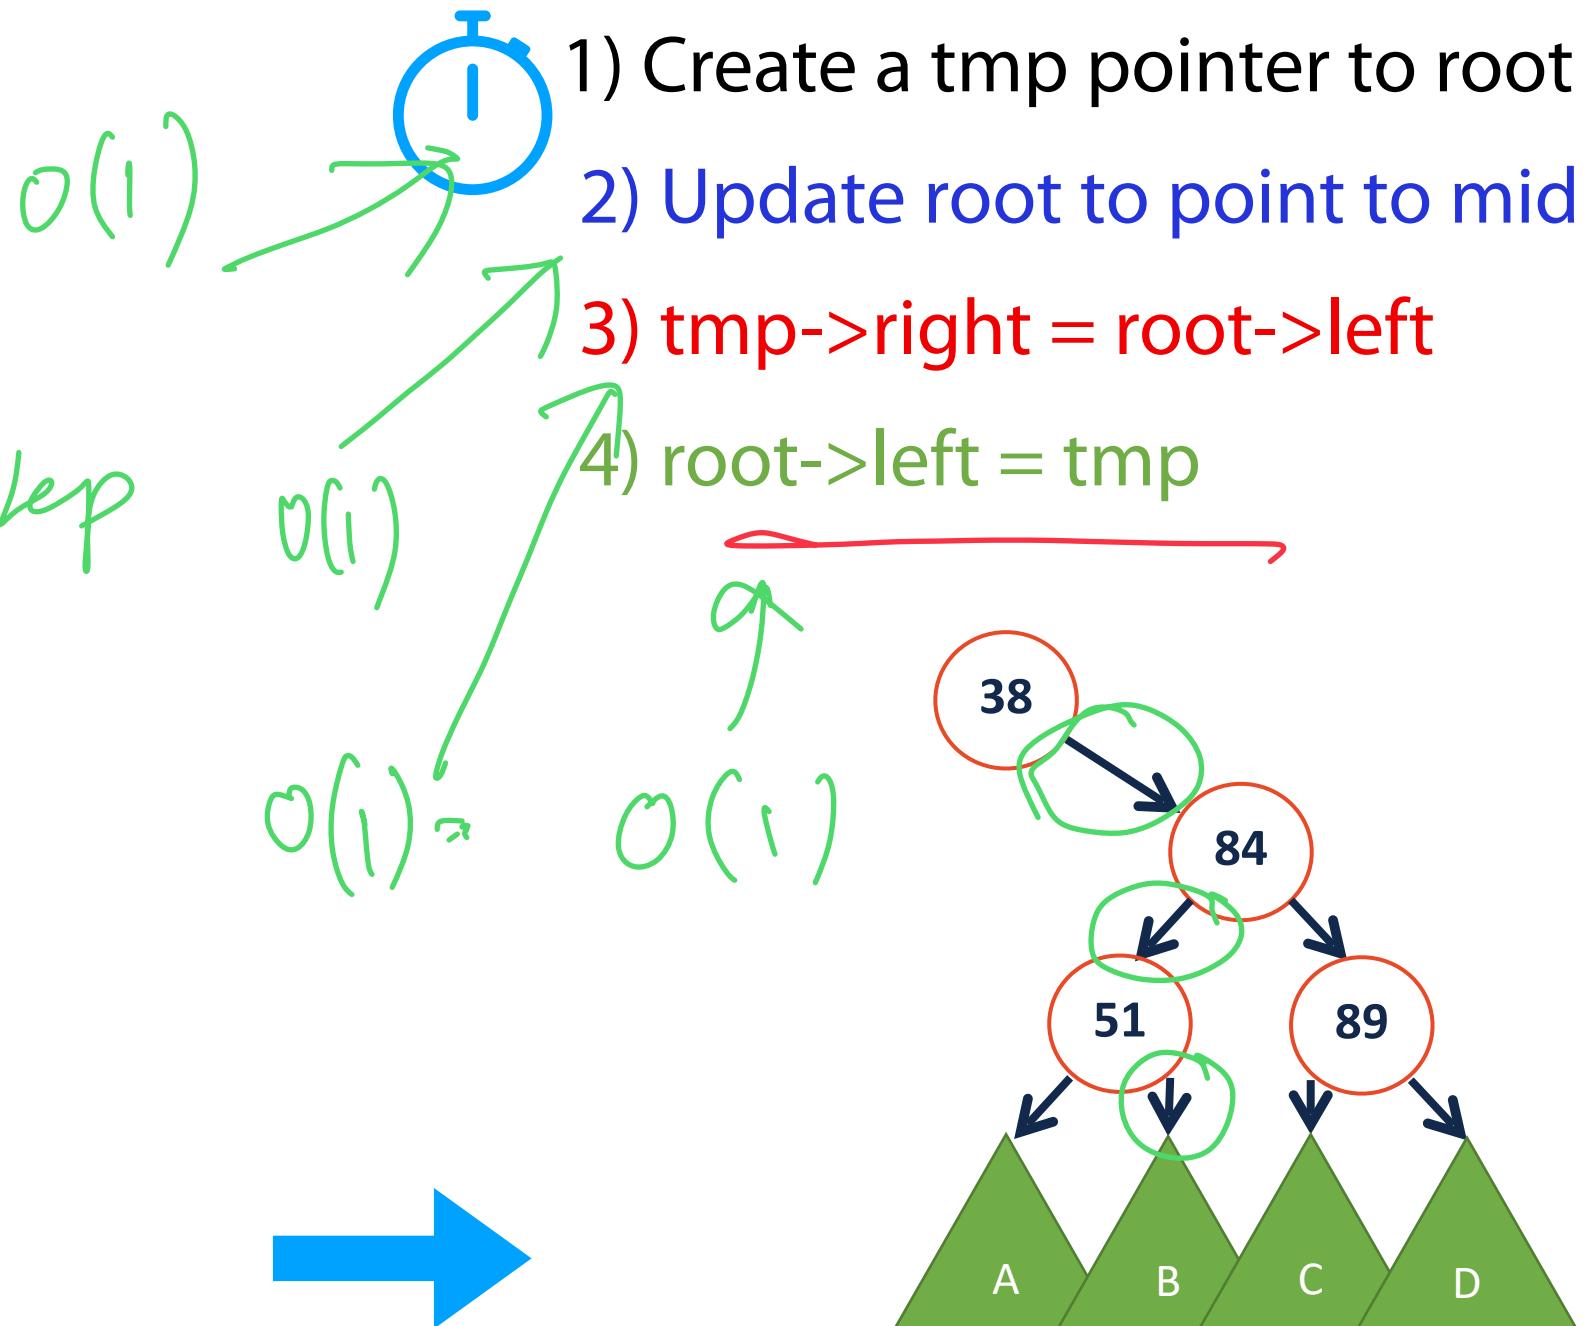
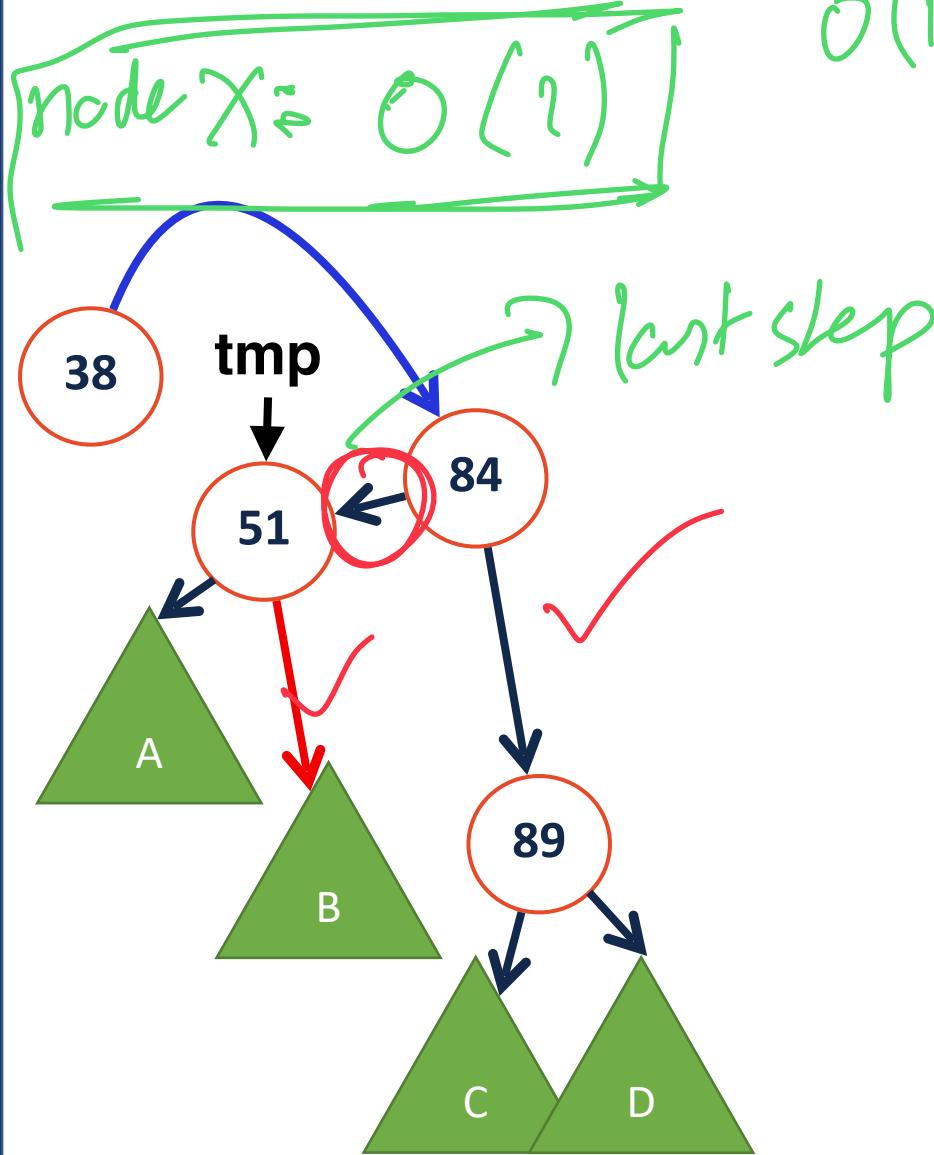
1) Create a tmp pointer to root

2) Update root to point to mid

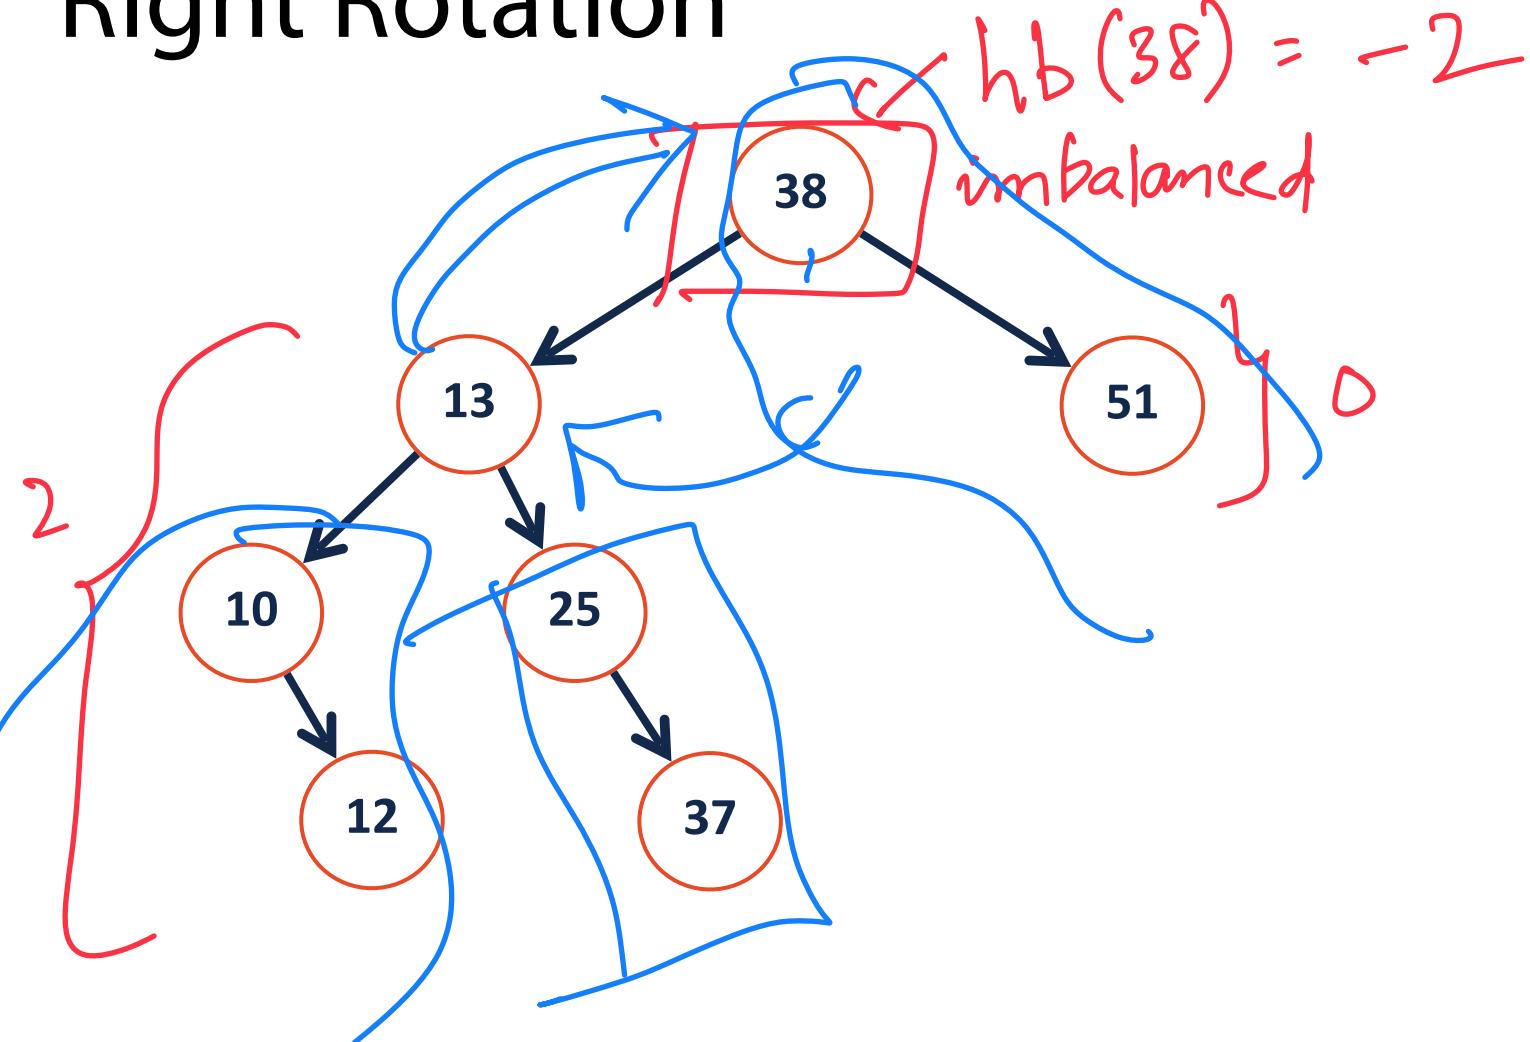
3) tmp->right = root->left



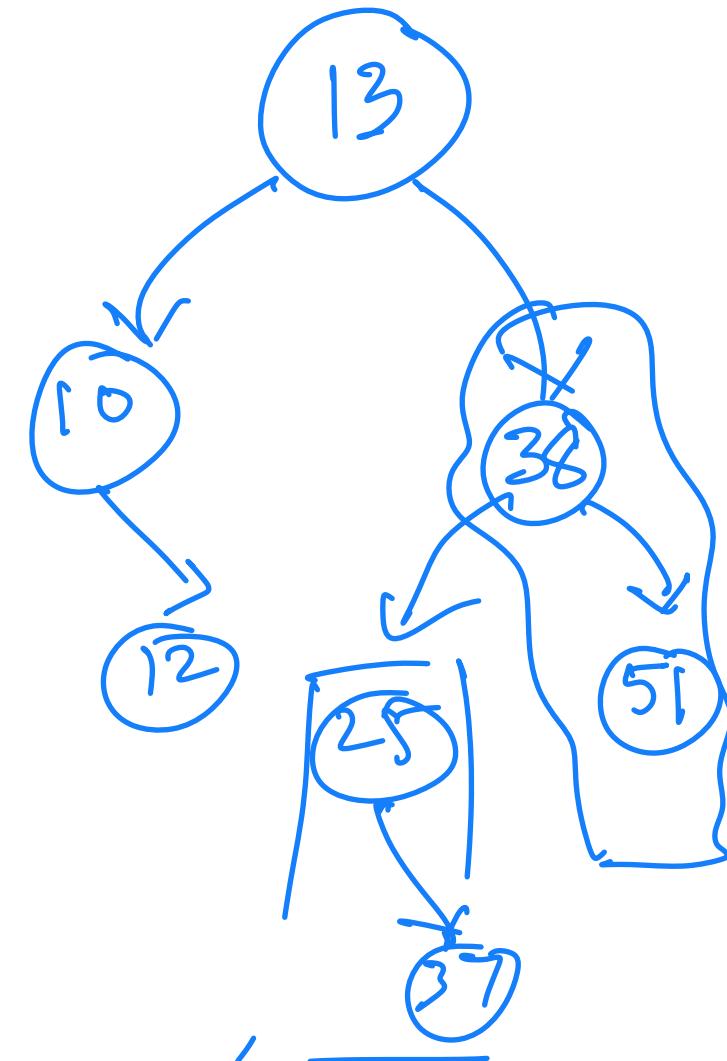
# Left Rotation



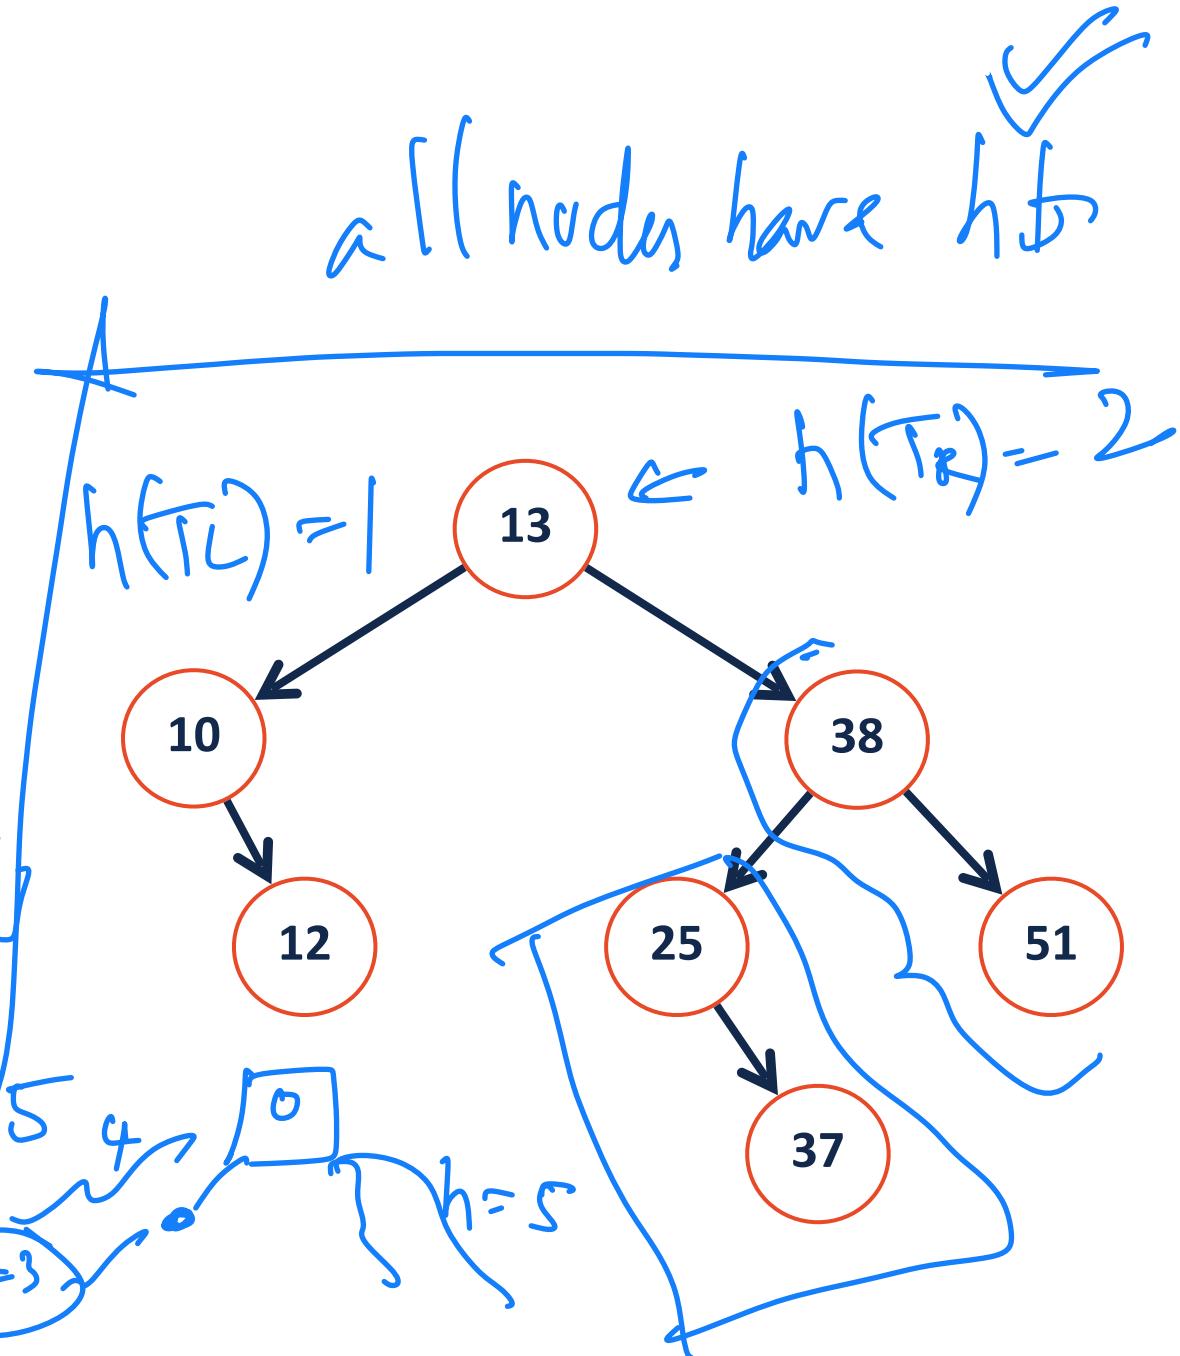
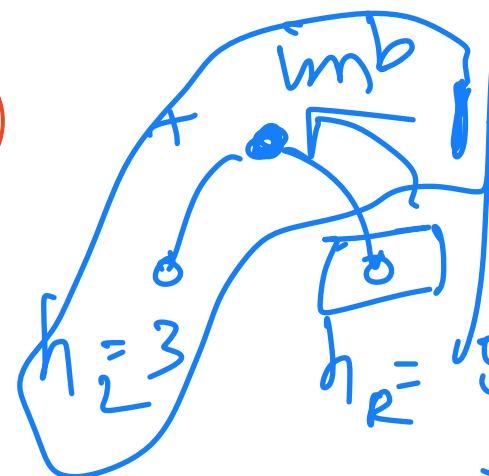
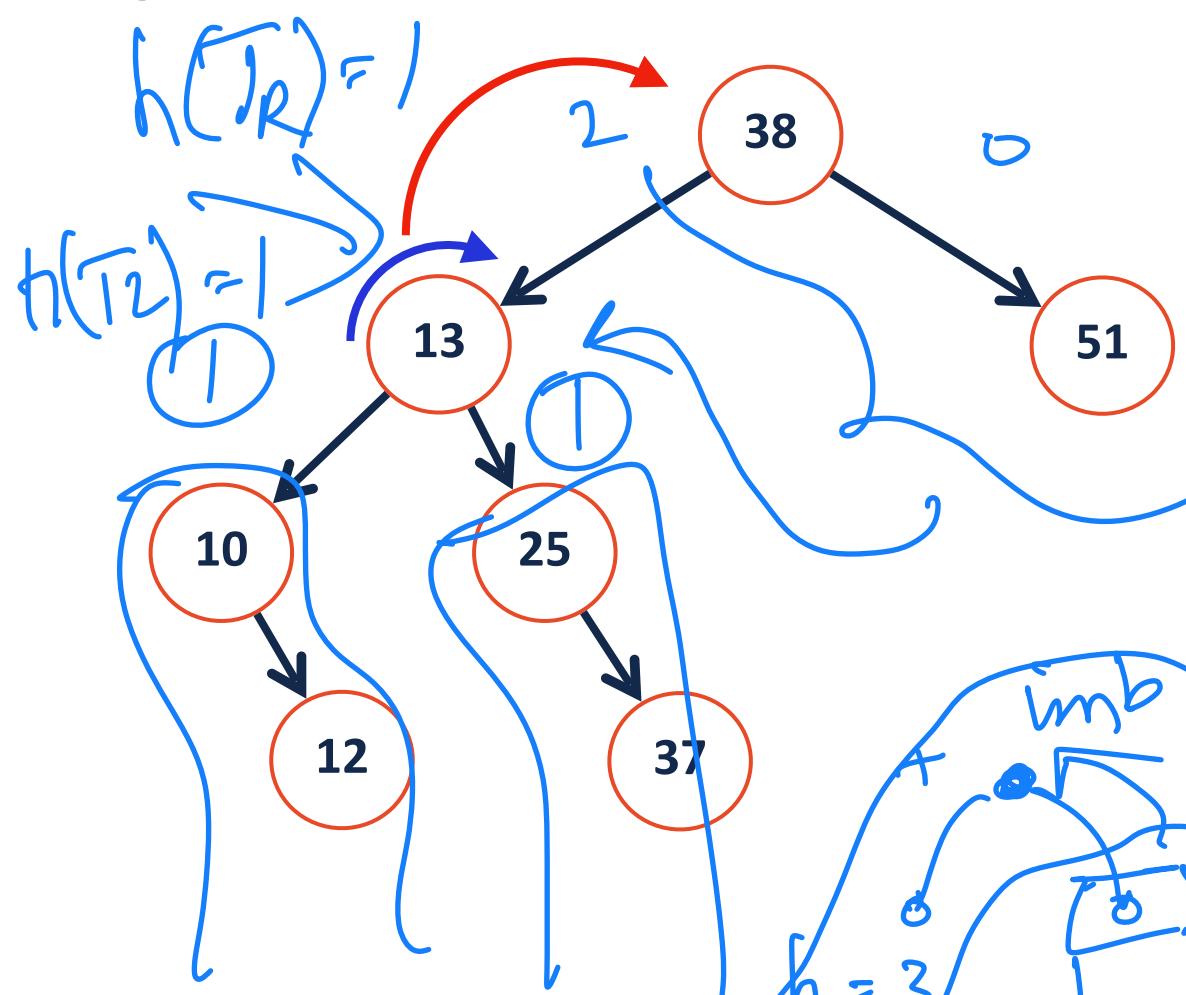
# Right Rotation



BSI



# Right Rotation

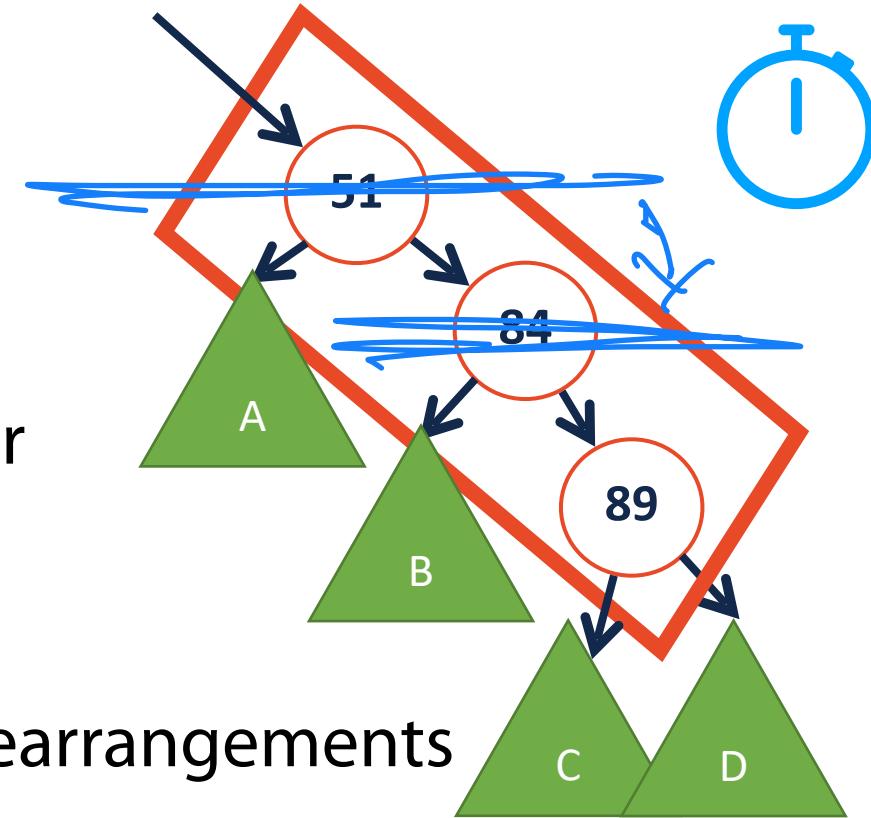


# Coding AVL Rotations

Two ways of visualizing:

1) Think of an arrow 'rotating' around the center

2) Recognize that there's a concrete order for rearrangements



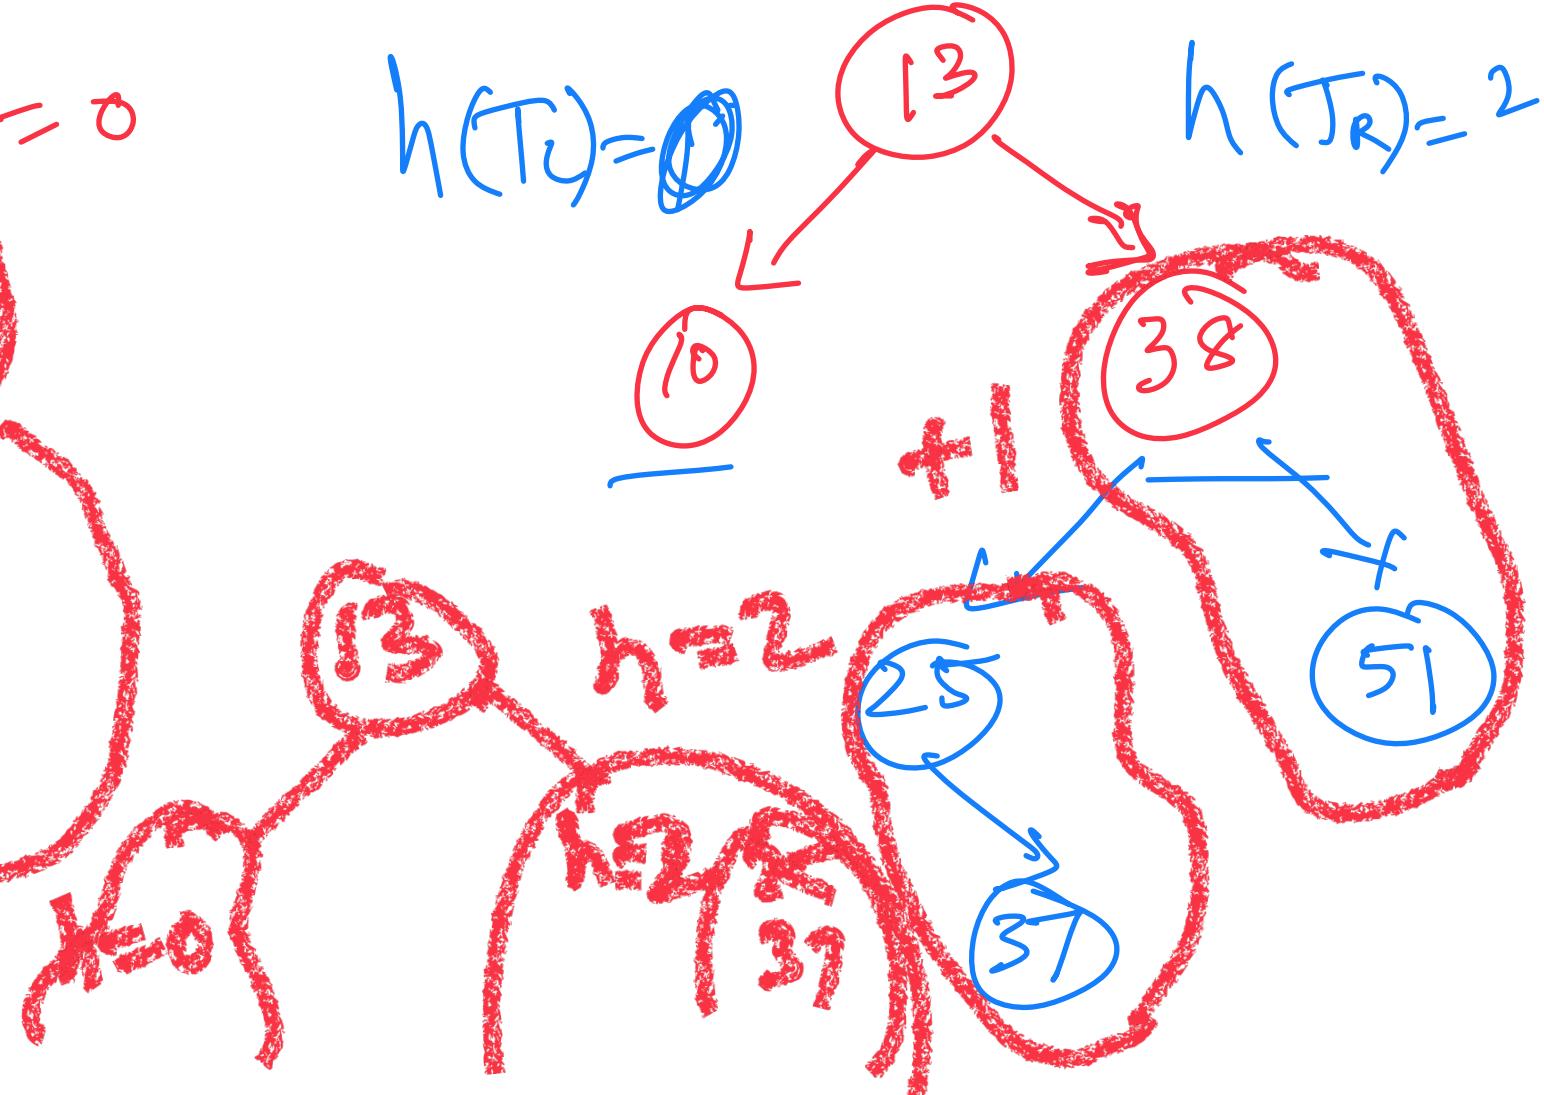
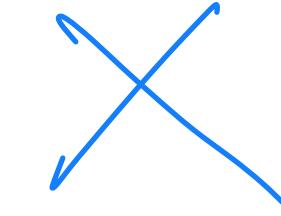
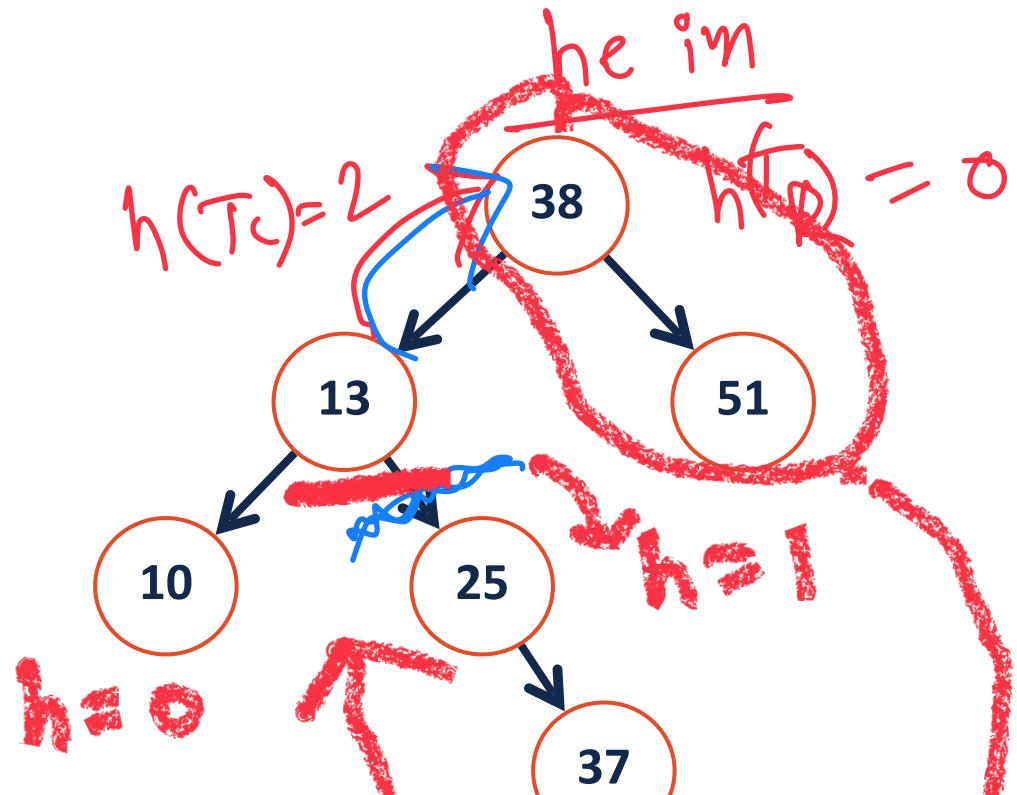
Ex: Unbalanced at current (root) node and need to *rotateLeft*?

Replace current (root) node with it's right child.

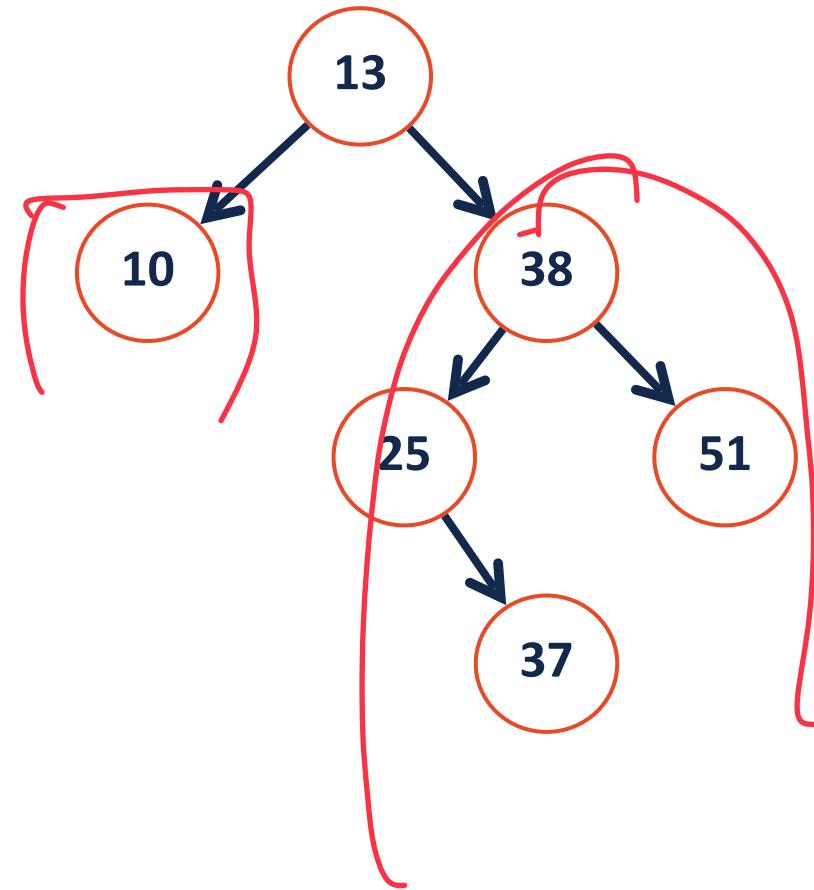
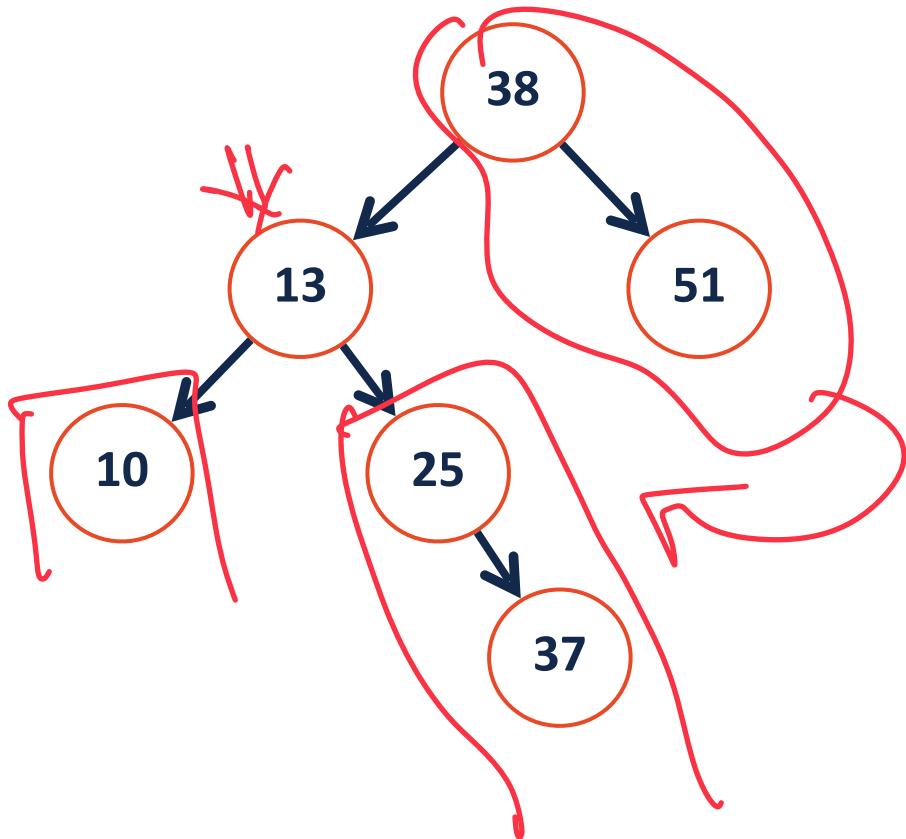
Set the right child's left child to be the current node's right

Make the current node the right child's left child

# AVL Rotation Practice

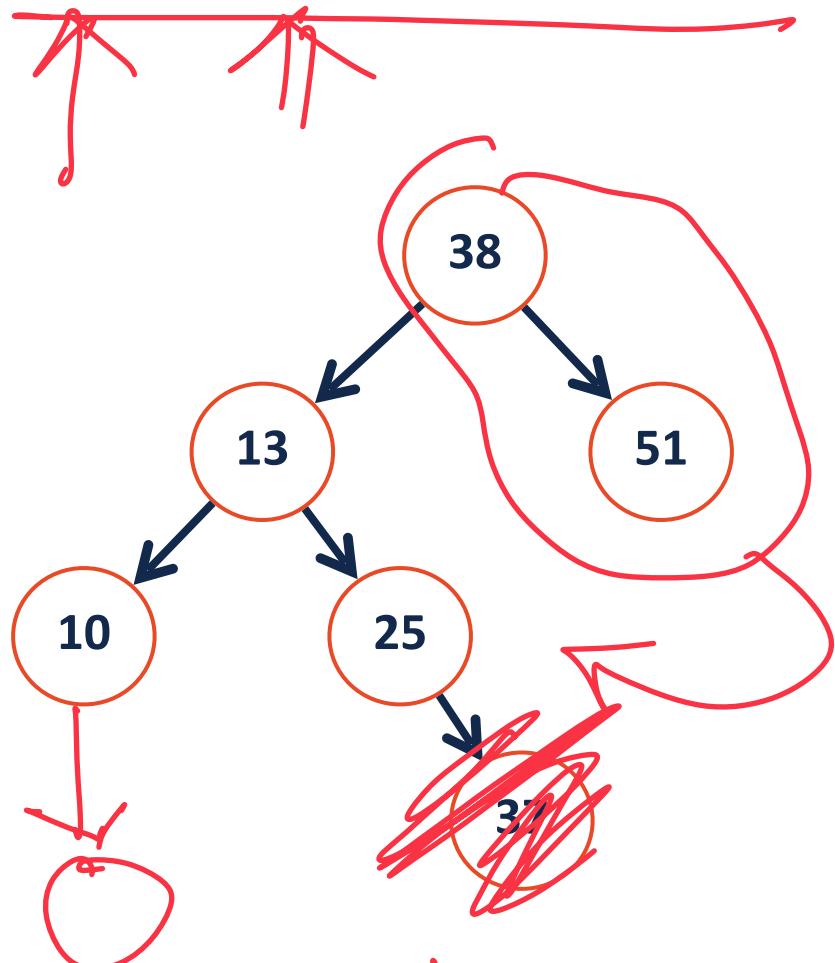


# AVL Rotation - Problems

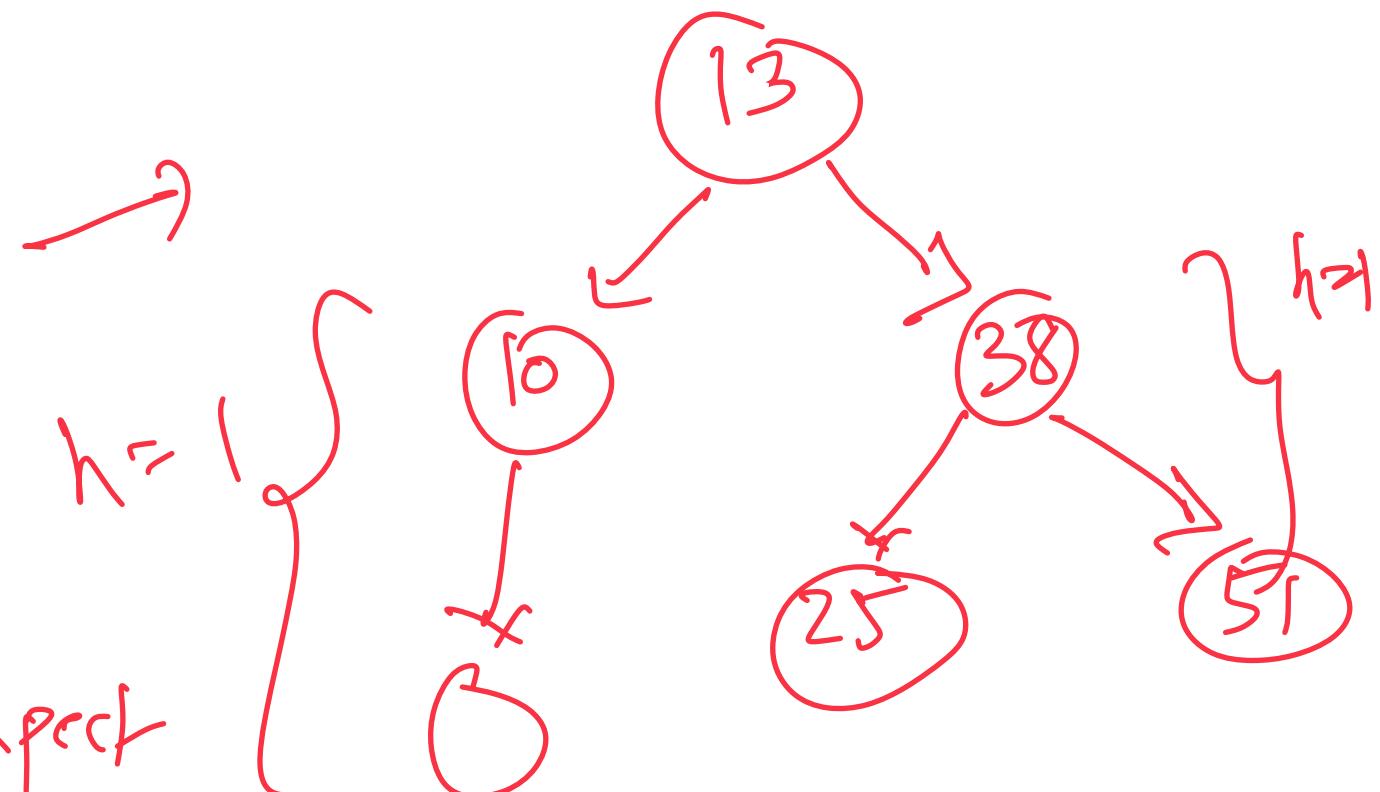


Something's not quite right...

# LeftRight Rotation

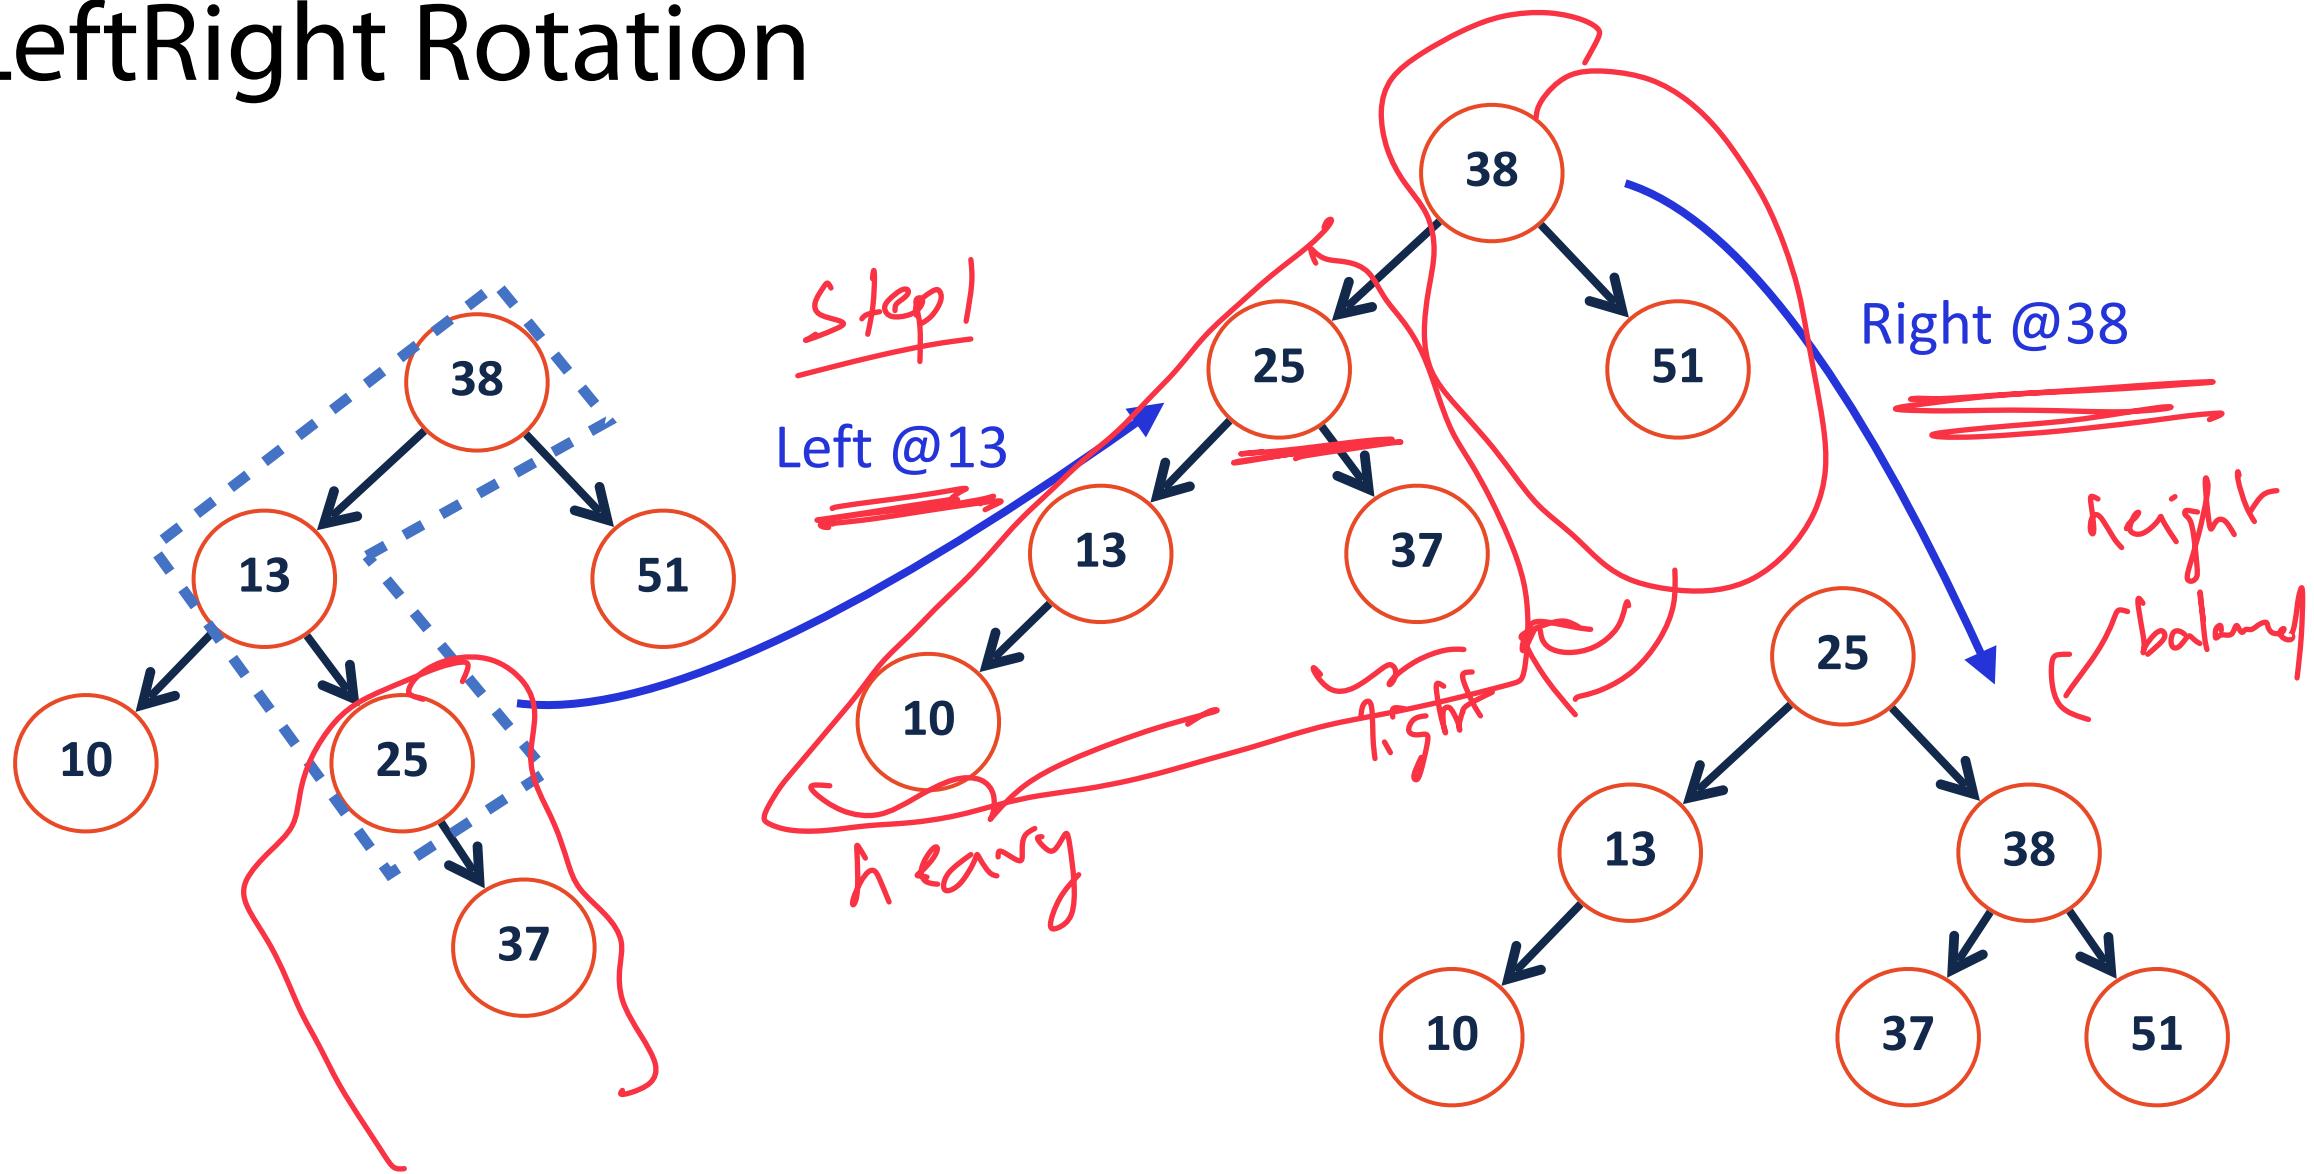


right heavy left imbalance

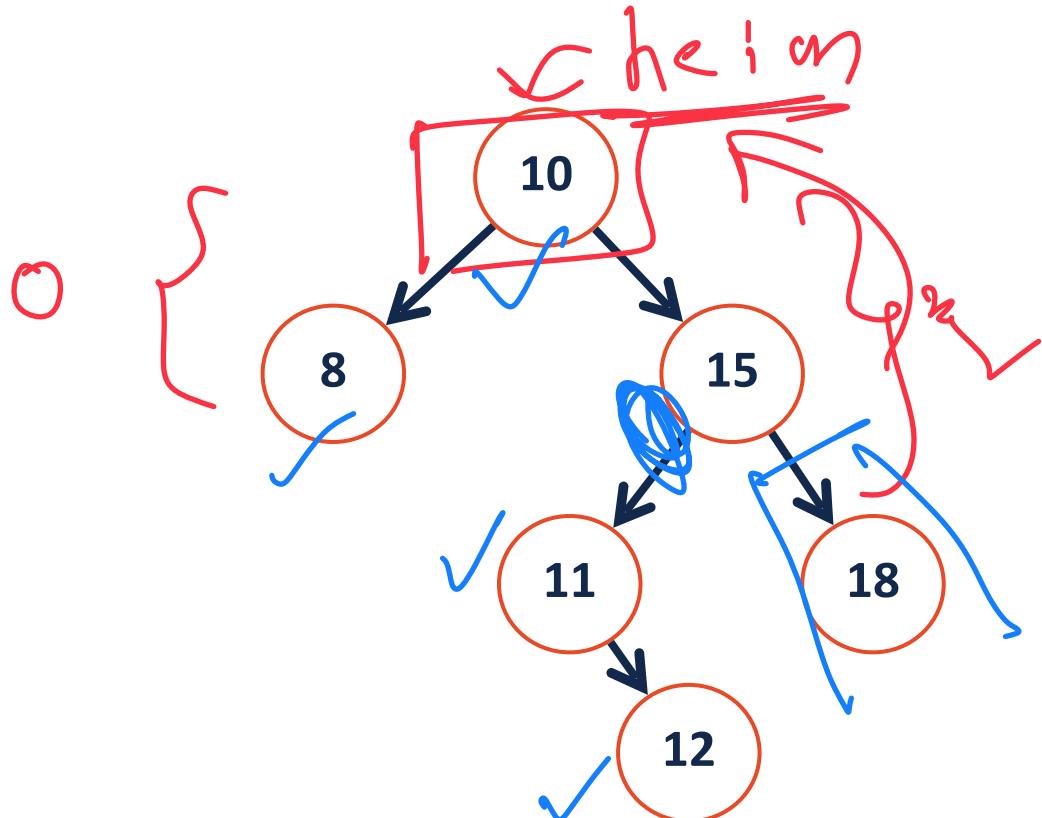


point: because we expect  
right of 13 to be overloaded, make it left heavy.

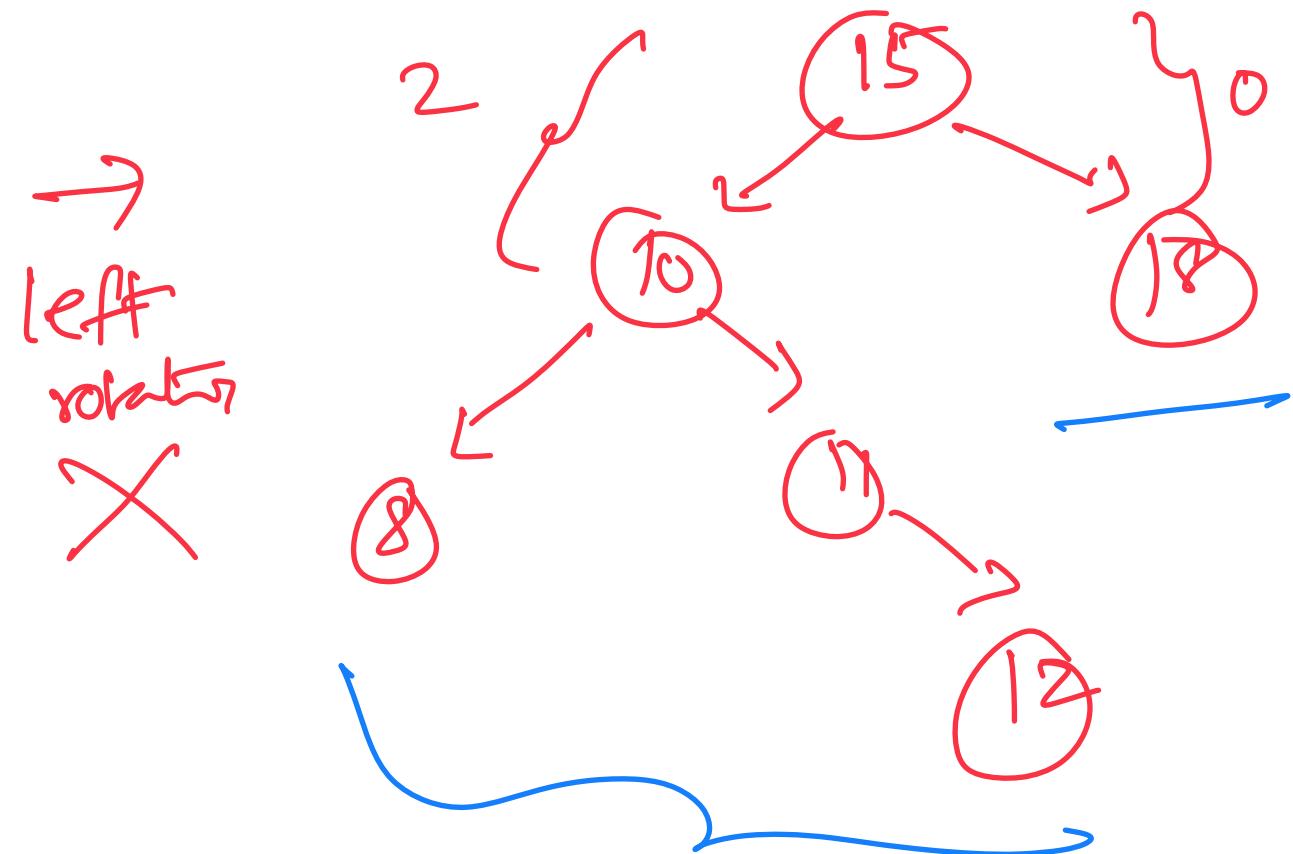
# LeftRight Rotation



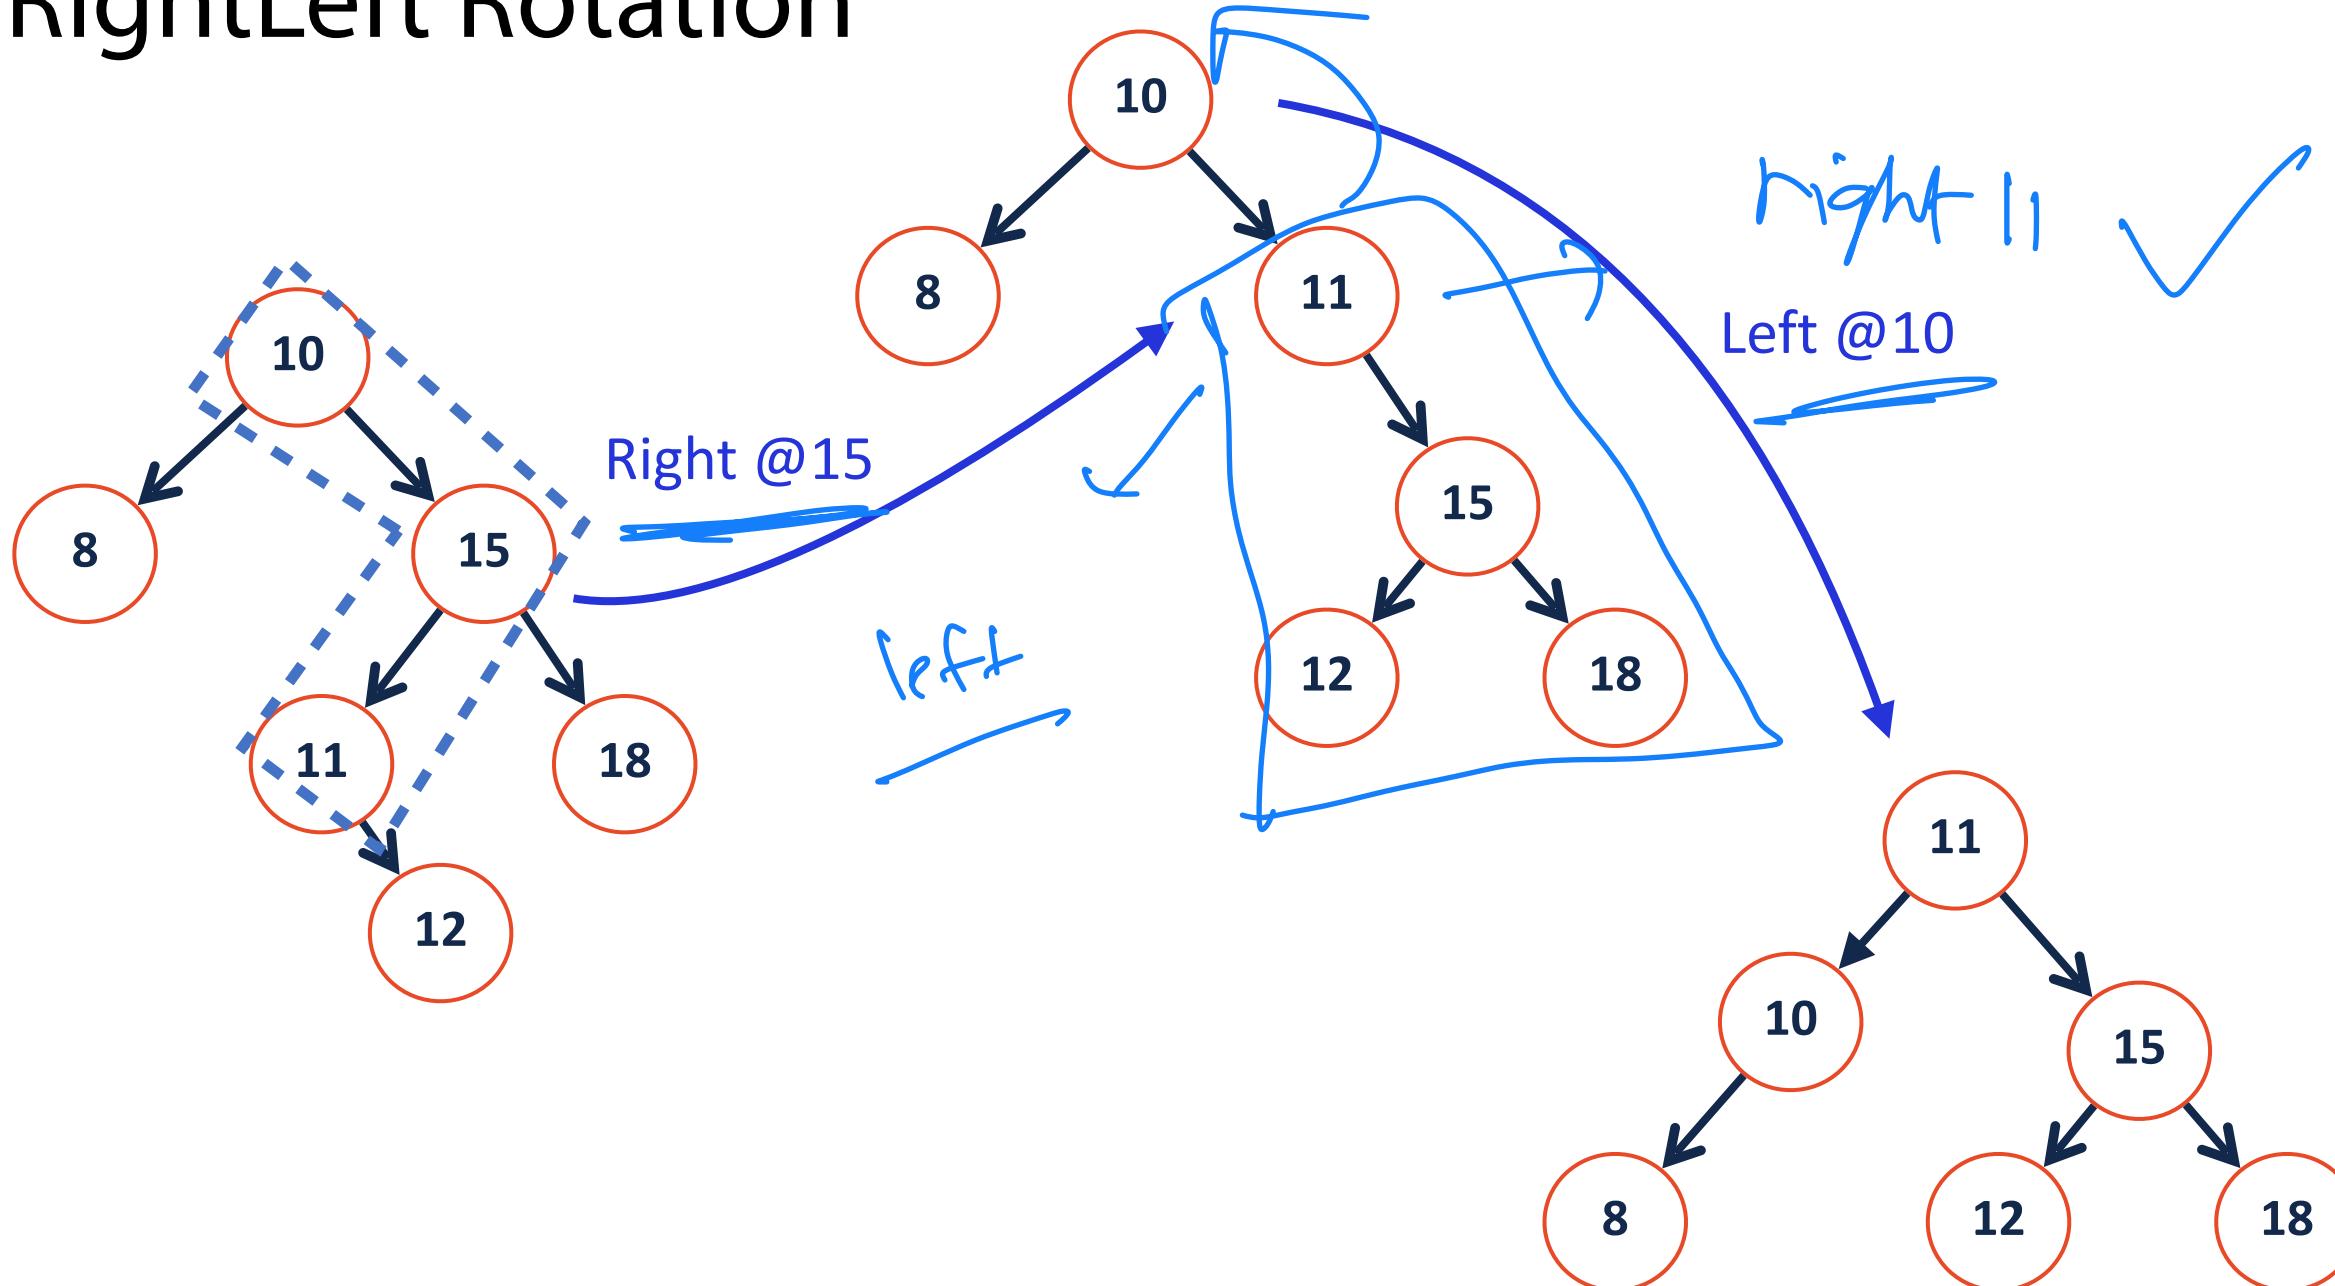
# RightLeft Rotation



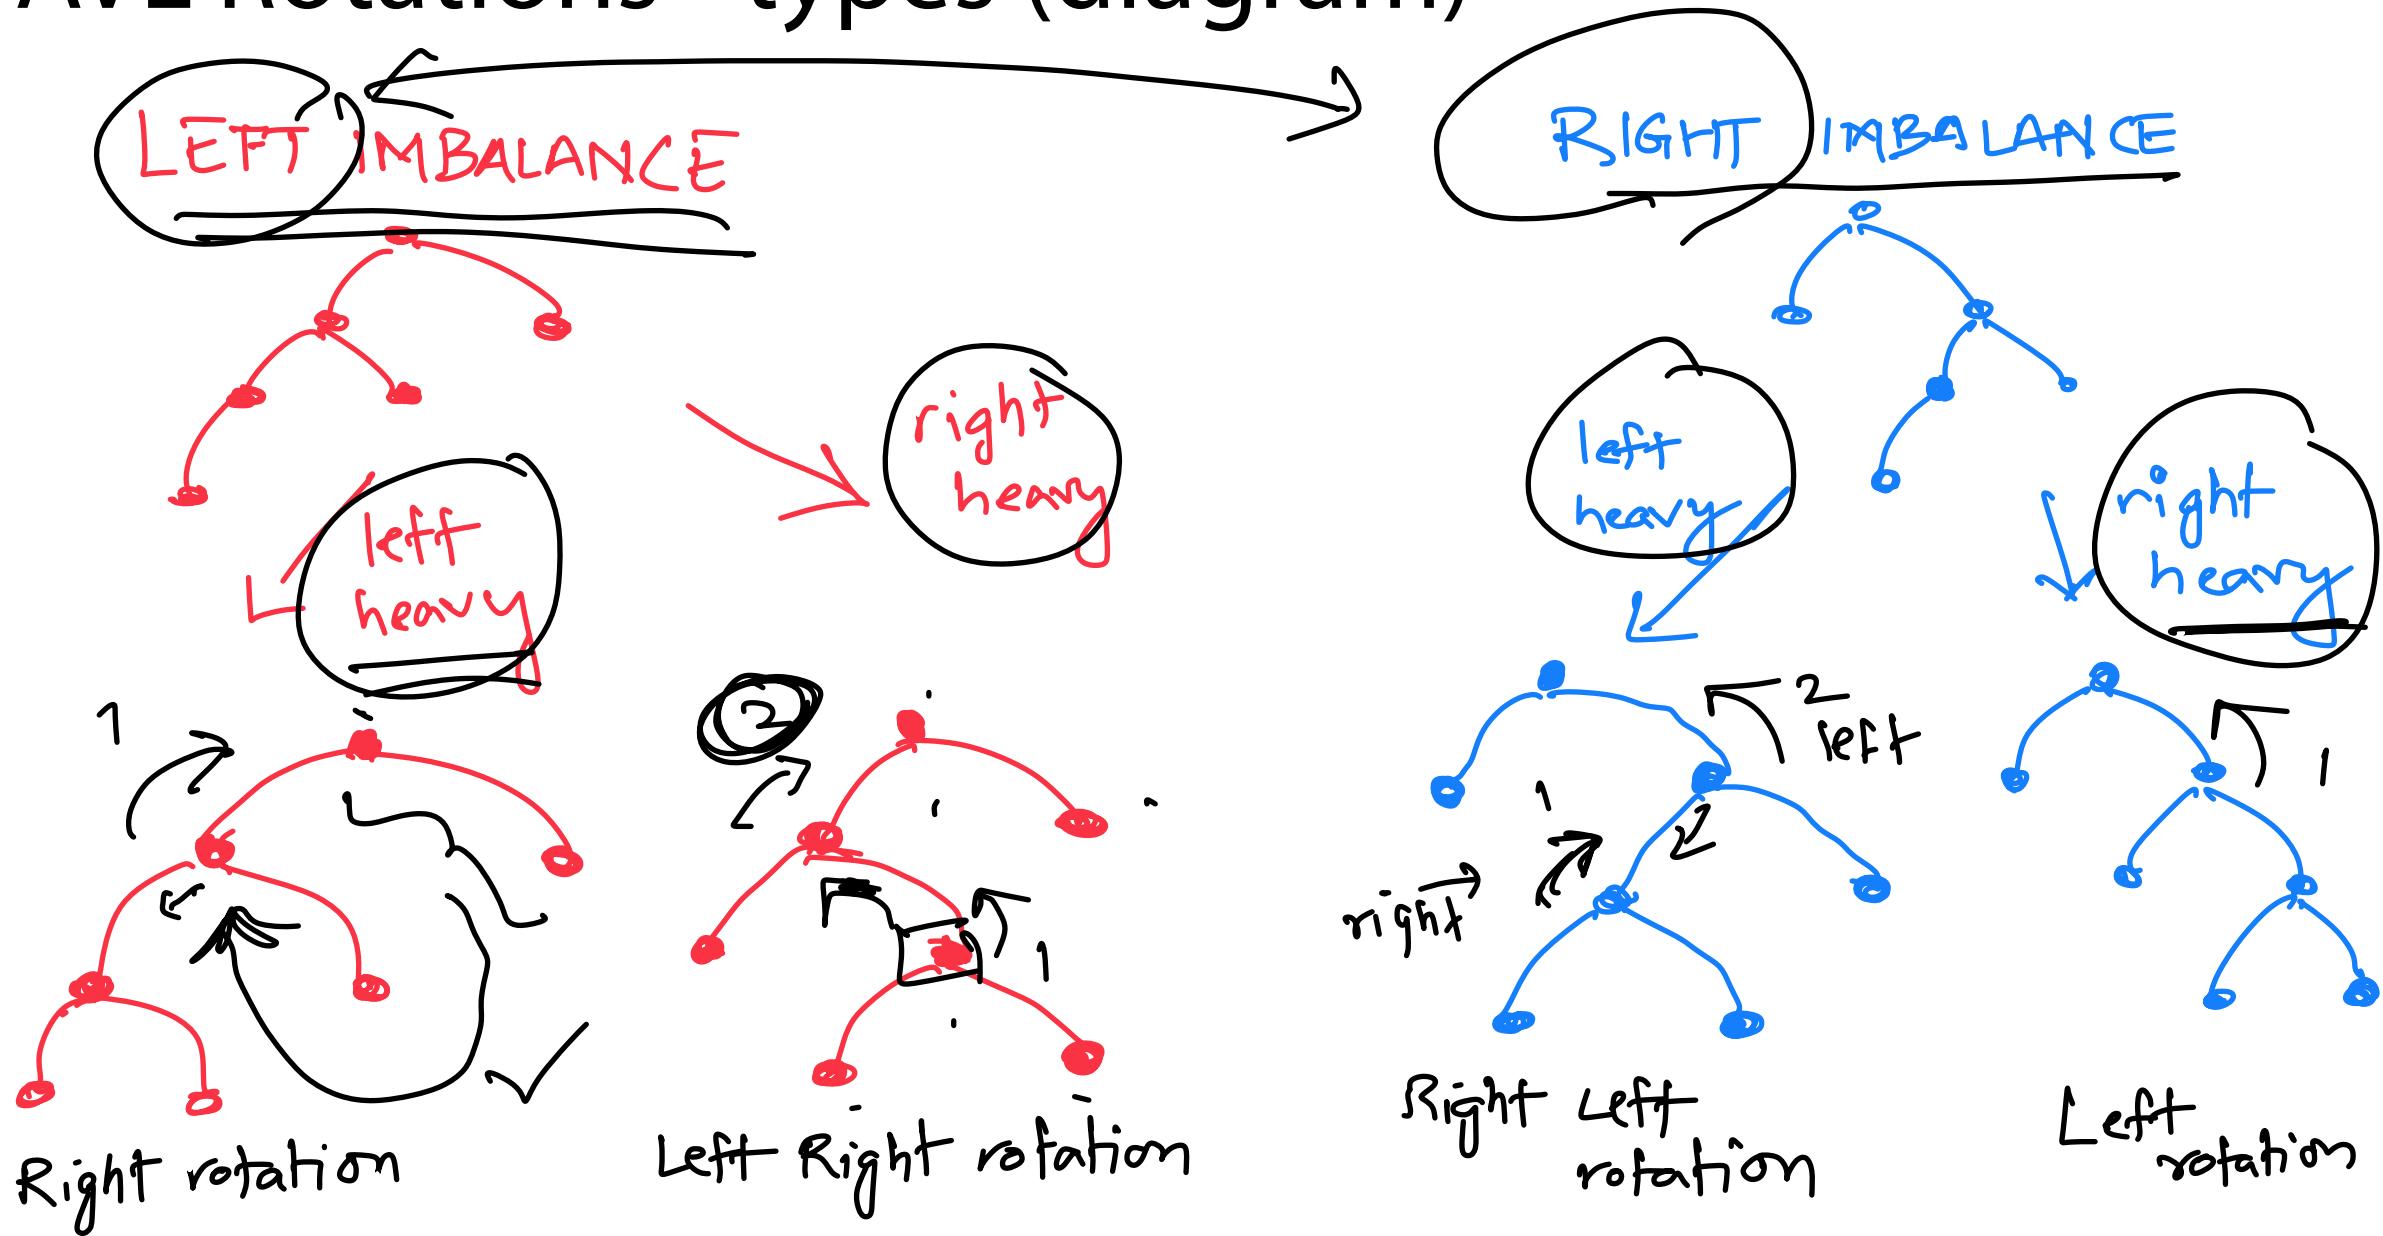
Left heavy right imbalance



# RightLeft Rotation



# AVL Rotations - types (diagram)



# AVL Rotations - types

1. Right Rotation

3. Left-Right Rotation

2. Left Rotation

4. Right-Left Rotation

# AVL Rotations



Four kinds of rotations: (L, R, LR, RL)

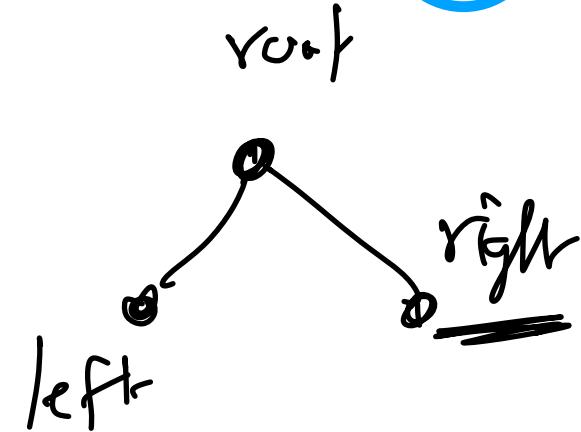
1. All rotations are local (subtrees are not impacted)

2. The running time of rotations are constant

3. The rotations maintain BST property

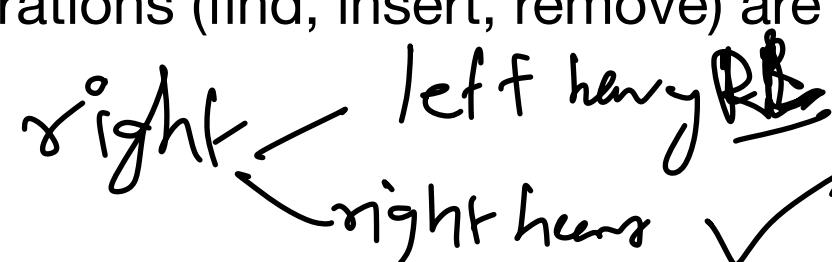
**Goal:**

Produce trees of height =  $O(\log n)$  so that all our major operations (find, insert, remove) are  $O(\log n)$  time.



$\text{right} \gg \text{left}$

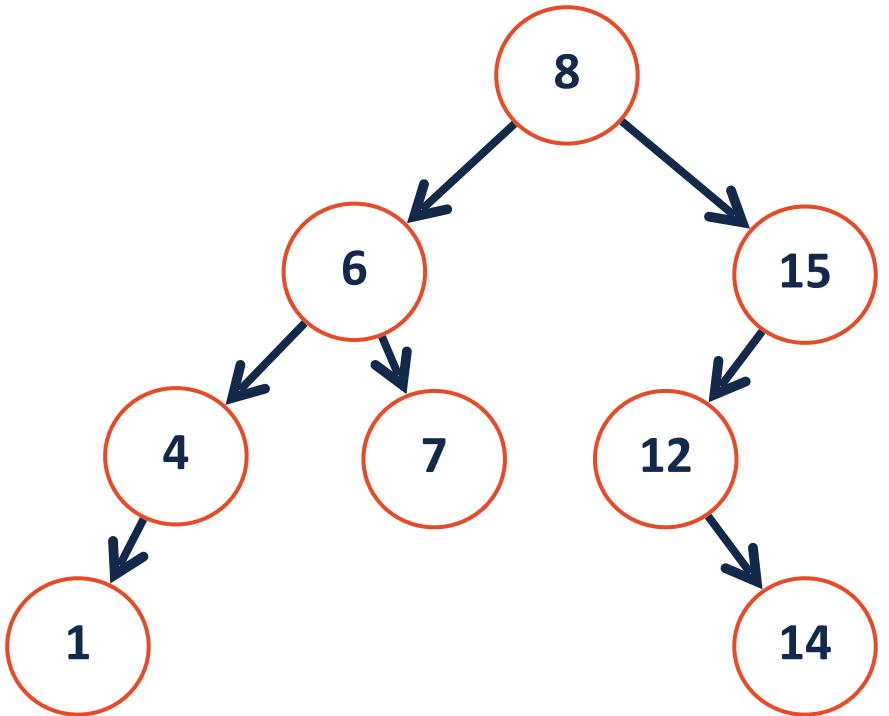
$\Rightarrow$  left rotation



# AVL Rotations - steps

1. Identify nodes with  $|height\ balance| \geq 2$
2. In a bottom up manner, fix nodes with  $|hb| \geq 2$
3. Identify the type of rotation to apply - by considering heights and balance of (heavier) child.
4. Execute rotation and return to parent in a bottom up manner until the entire tree is balanced.

# AVL Rotation Practice



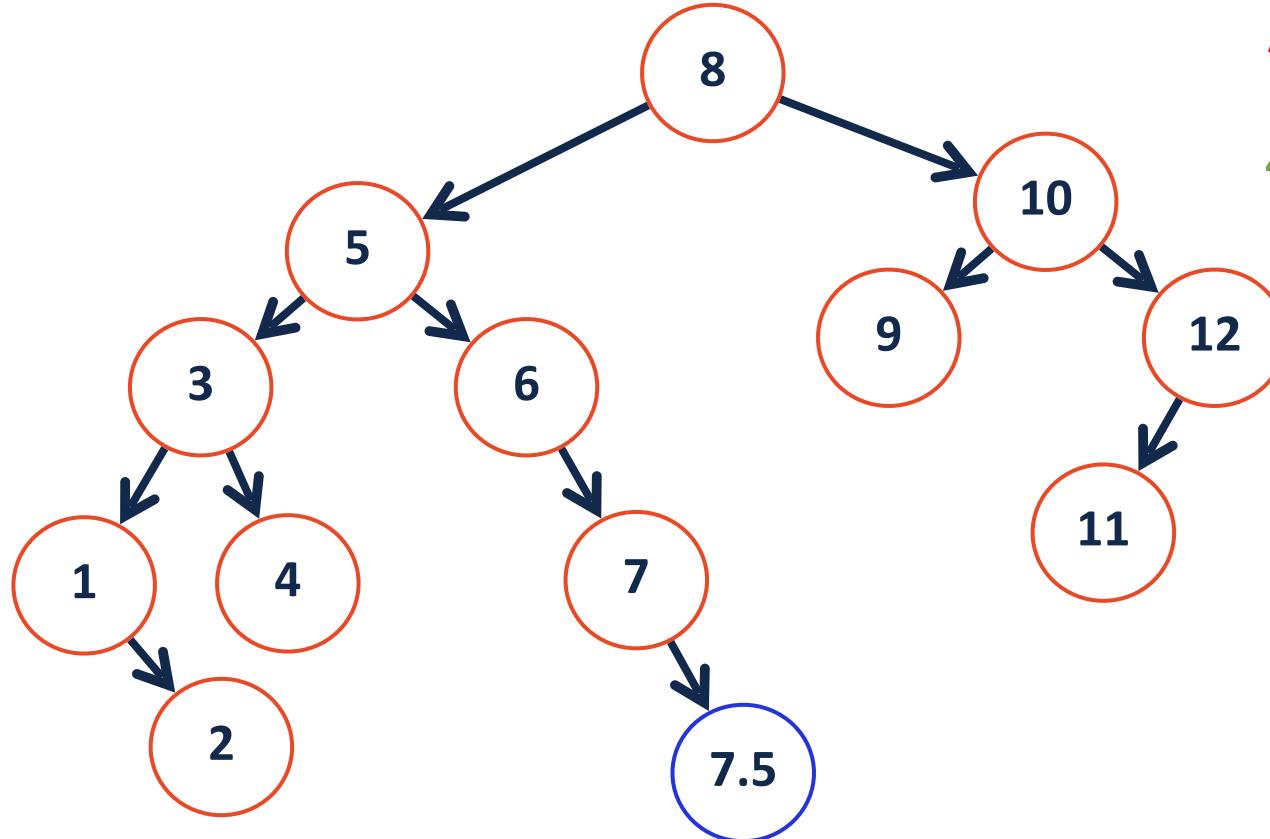
# AVL vs BST ADT

The AVL tree is a modified binary search tree that rotates **when necessary**

How does the constraint on balance affect the core functions?

Operation	BST $h = O(n)$	AVL tree $h = O(\log n)$	rotation
Find	$O(h)$	$O(\log n)$	
Insert	$O(h)$	$O(\log n)$	
Remove	$O(h)$	$O(\log n)$	

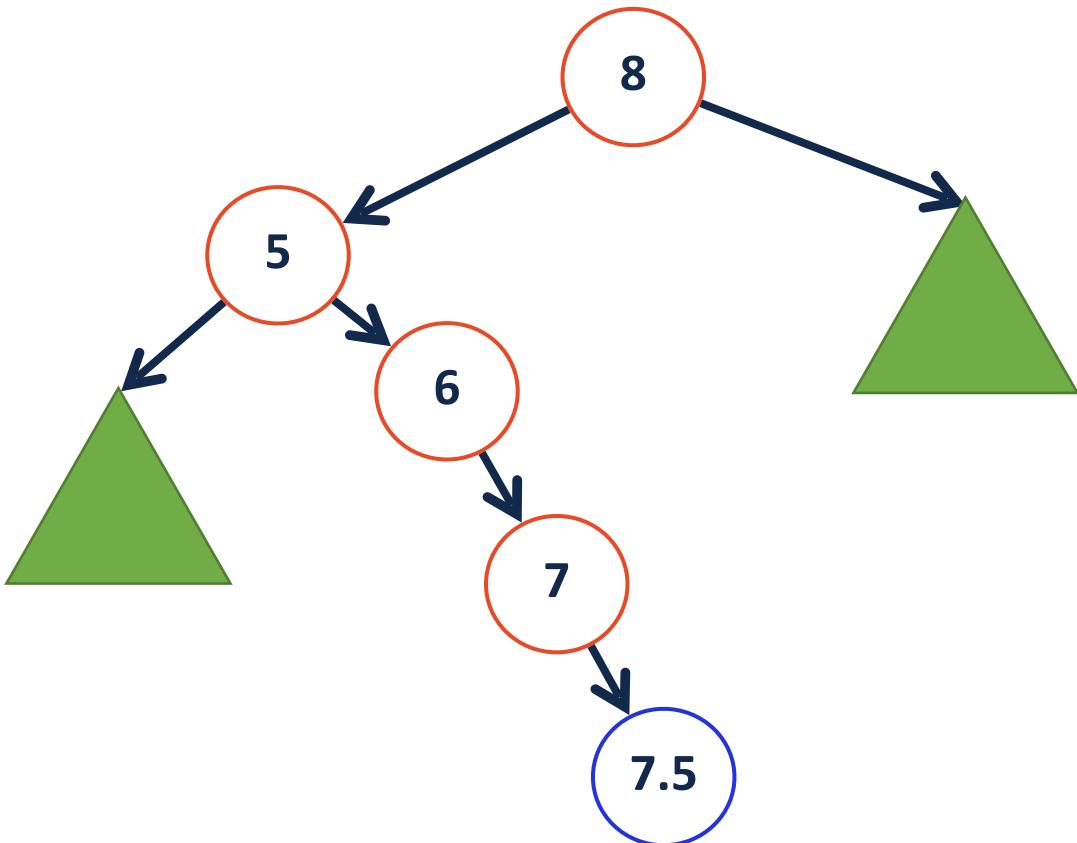
# Left Rotation



- 1) Create a tmp pointer to root
- 2) Update root to point to mid
- 3) `tmp->right = root->left`
- 4) `root->left = tmp`

# Left Rotation

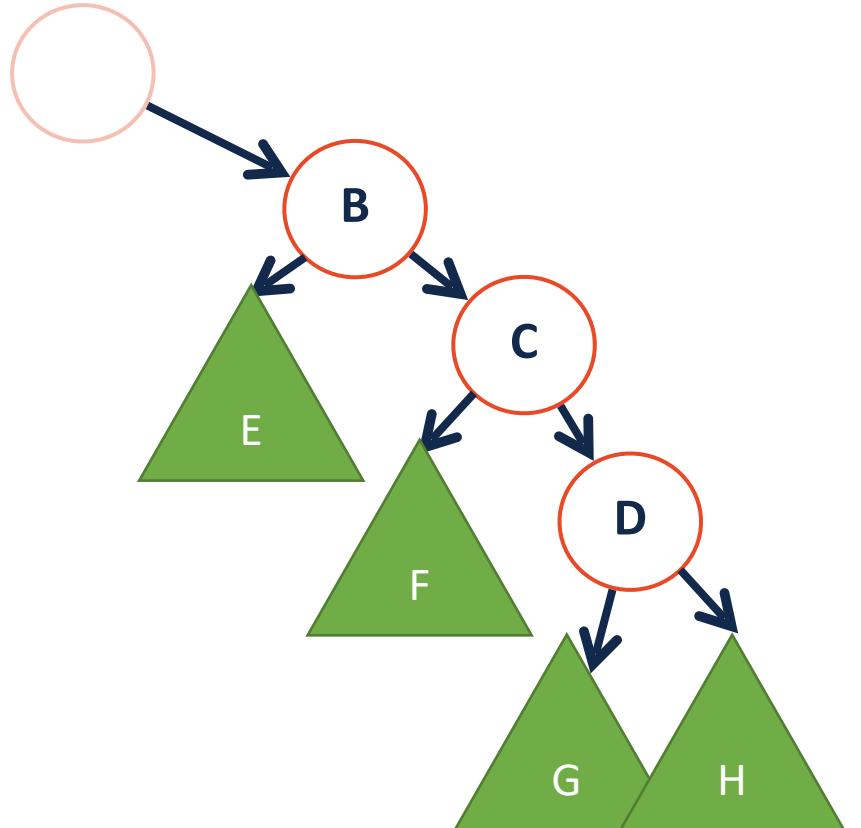
All rotations are local (subtrees are not impacted)



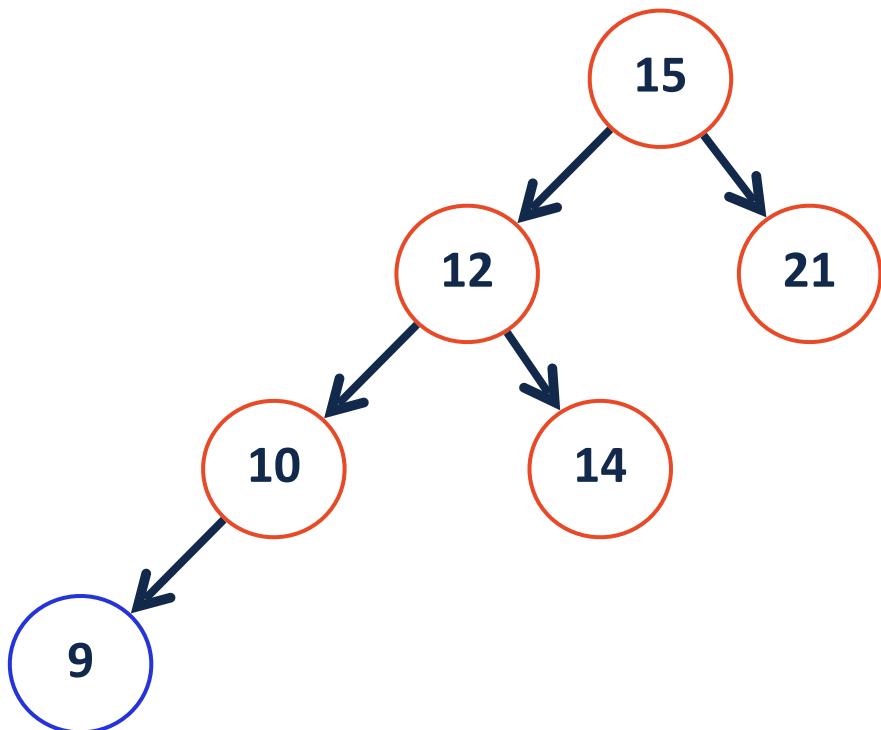
# Left Rotation



All rotations preserve BST property

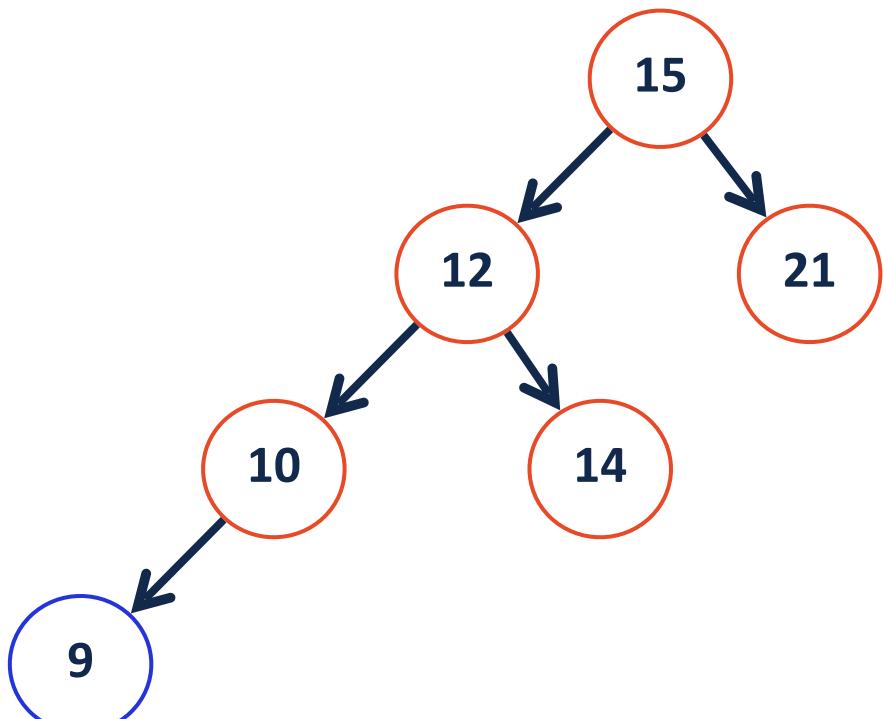


# Right Rotation

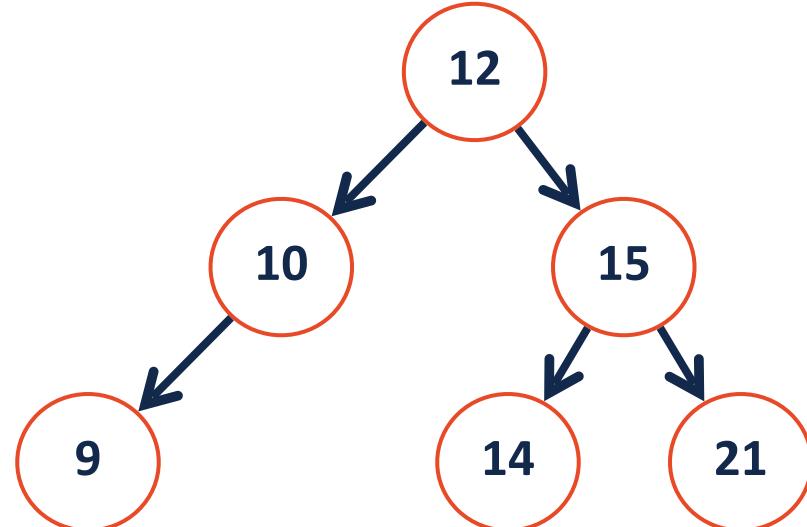


- 1) Create a tmp pointer to root
- 2) Update root to point to mid
- 3) **tmp->left = root->right**
- 4) **root->right = tmp**

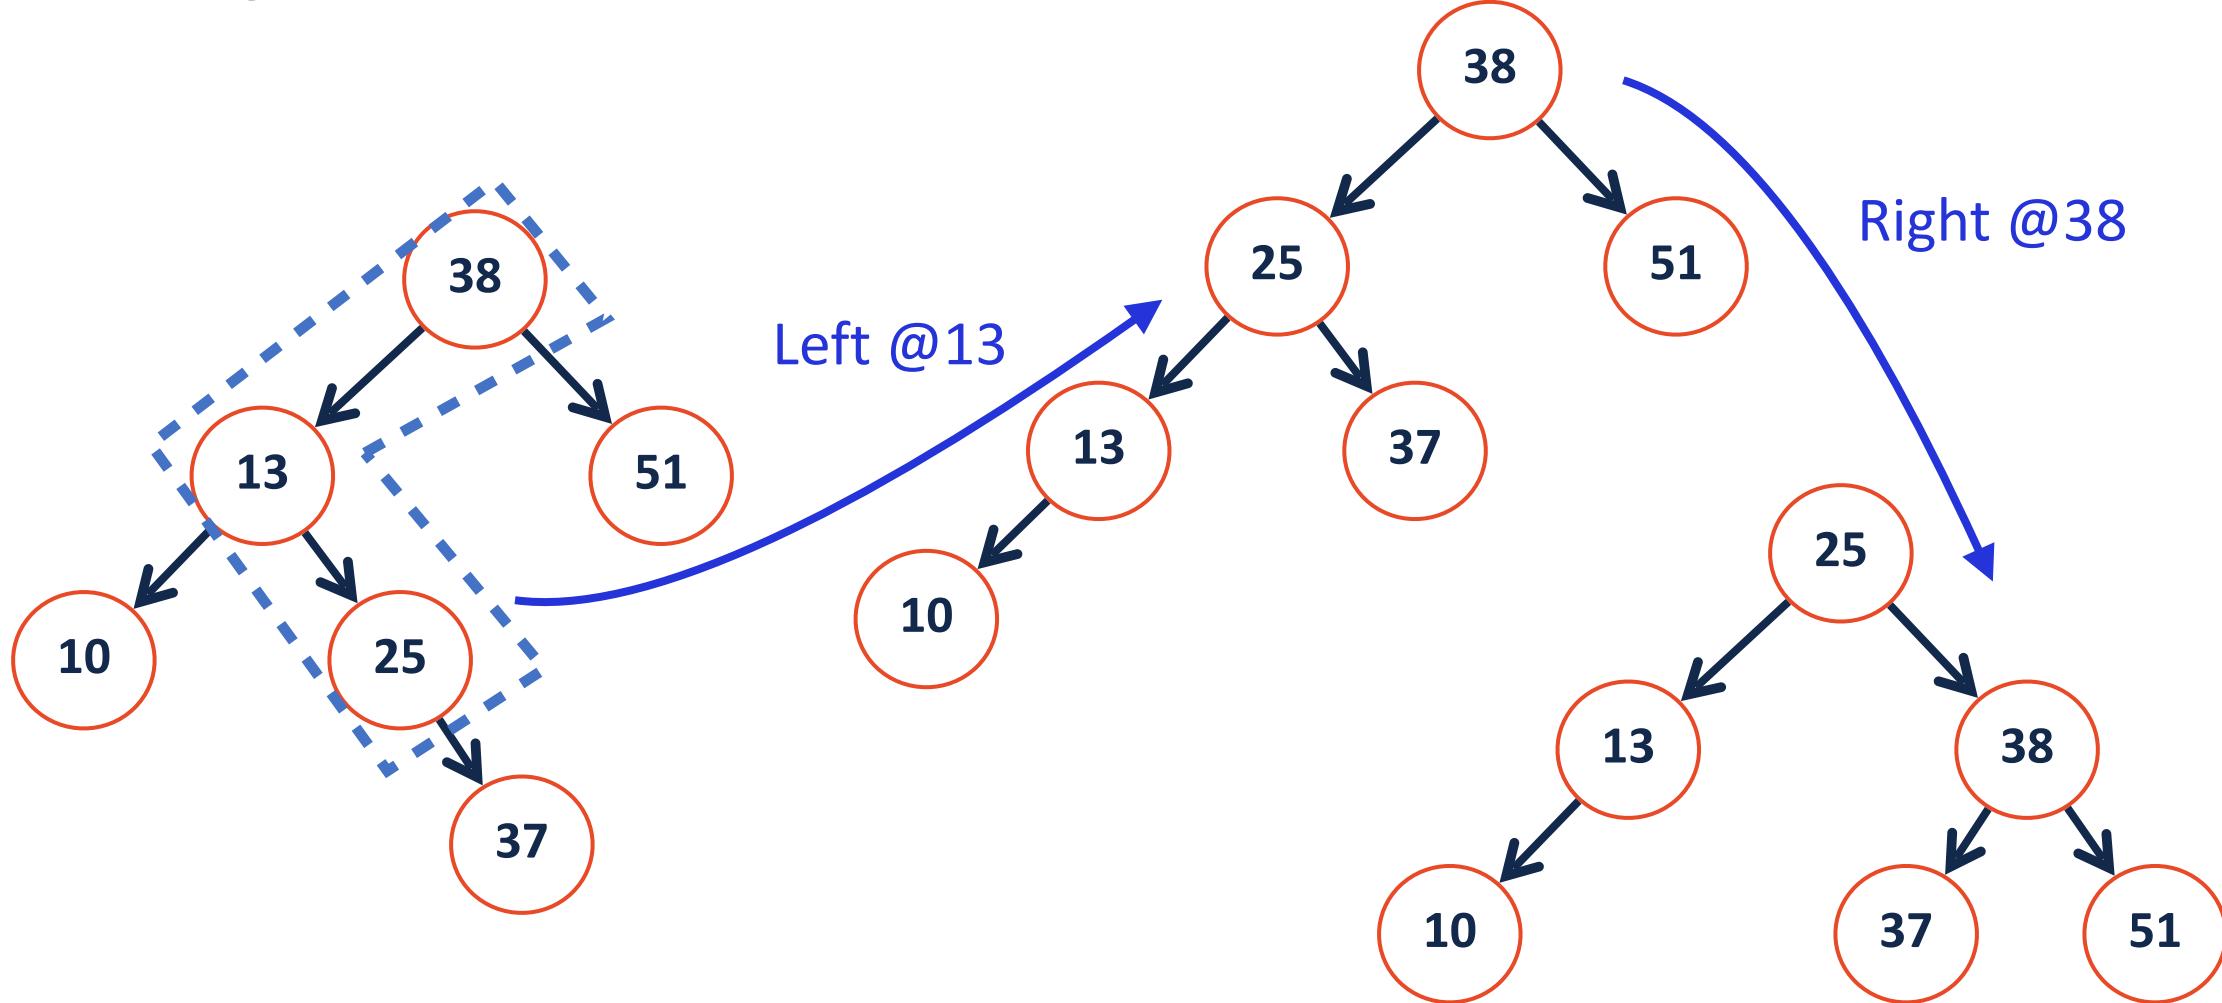
# Right Rotation



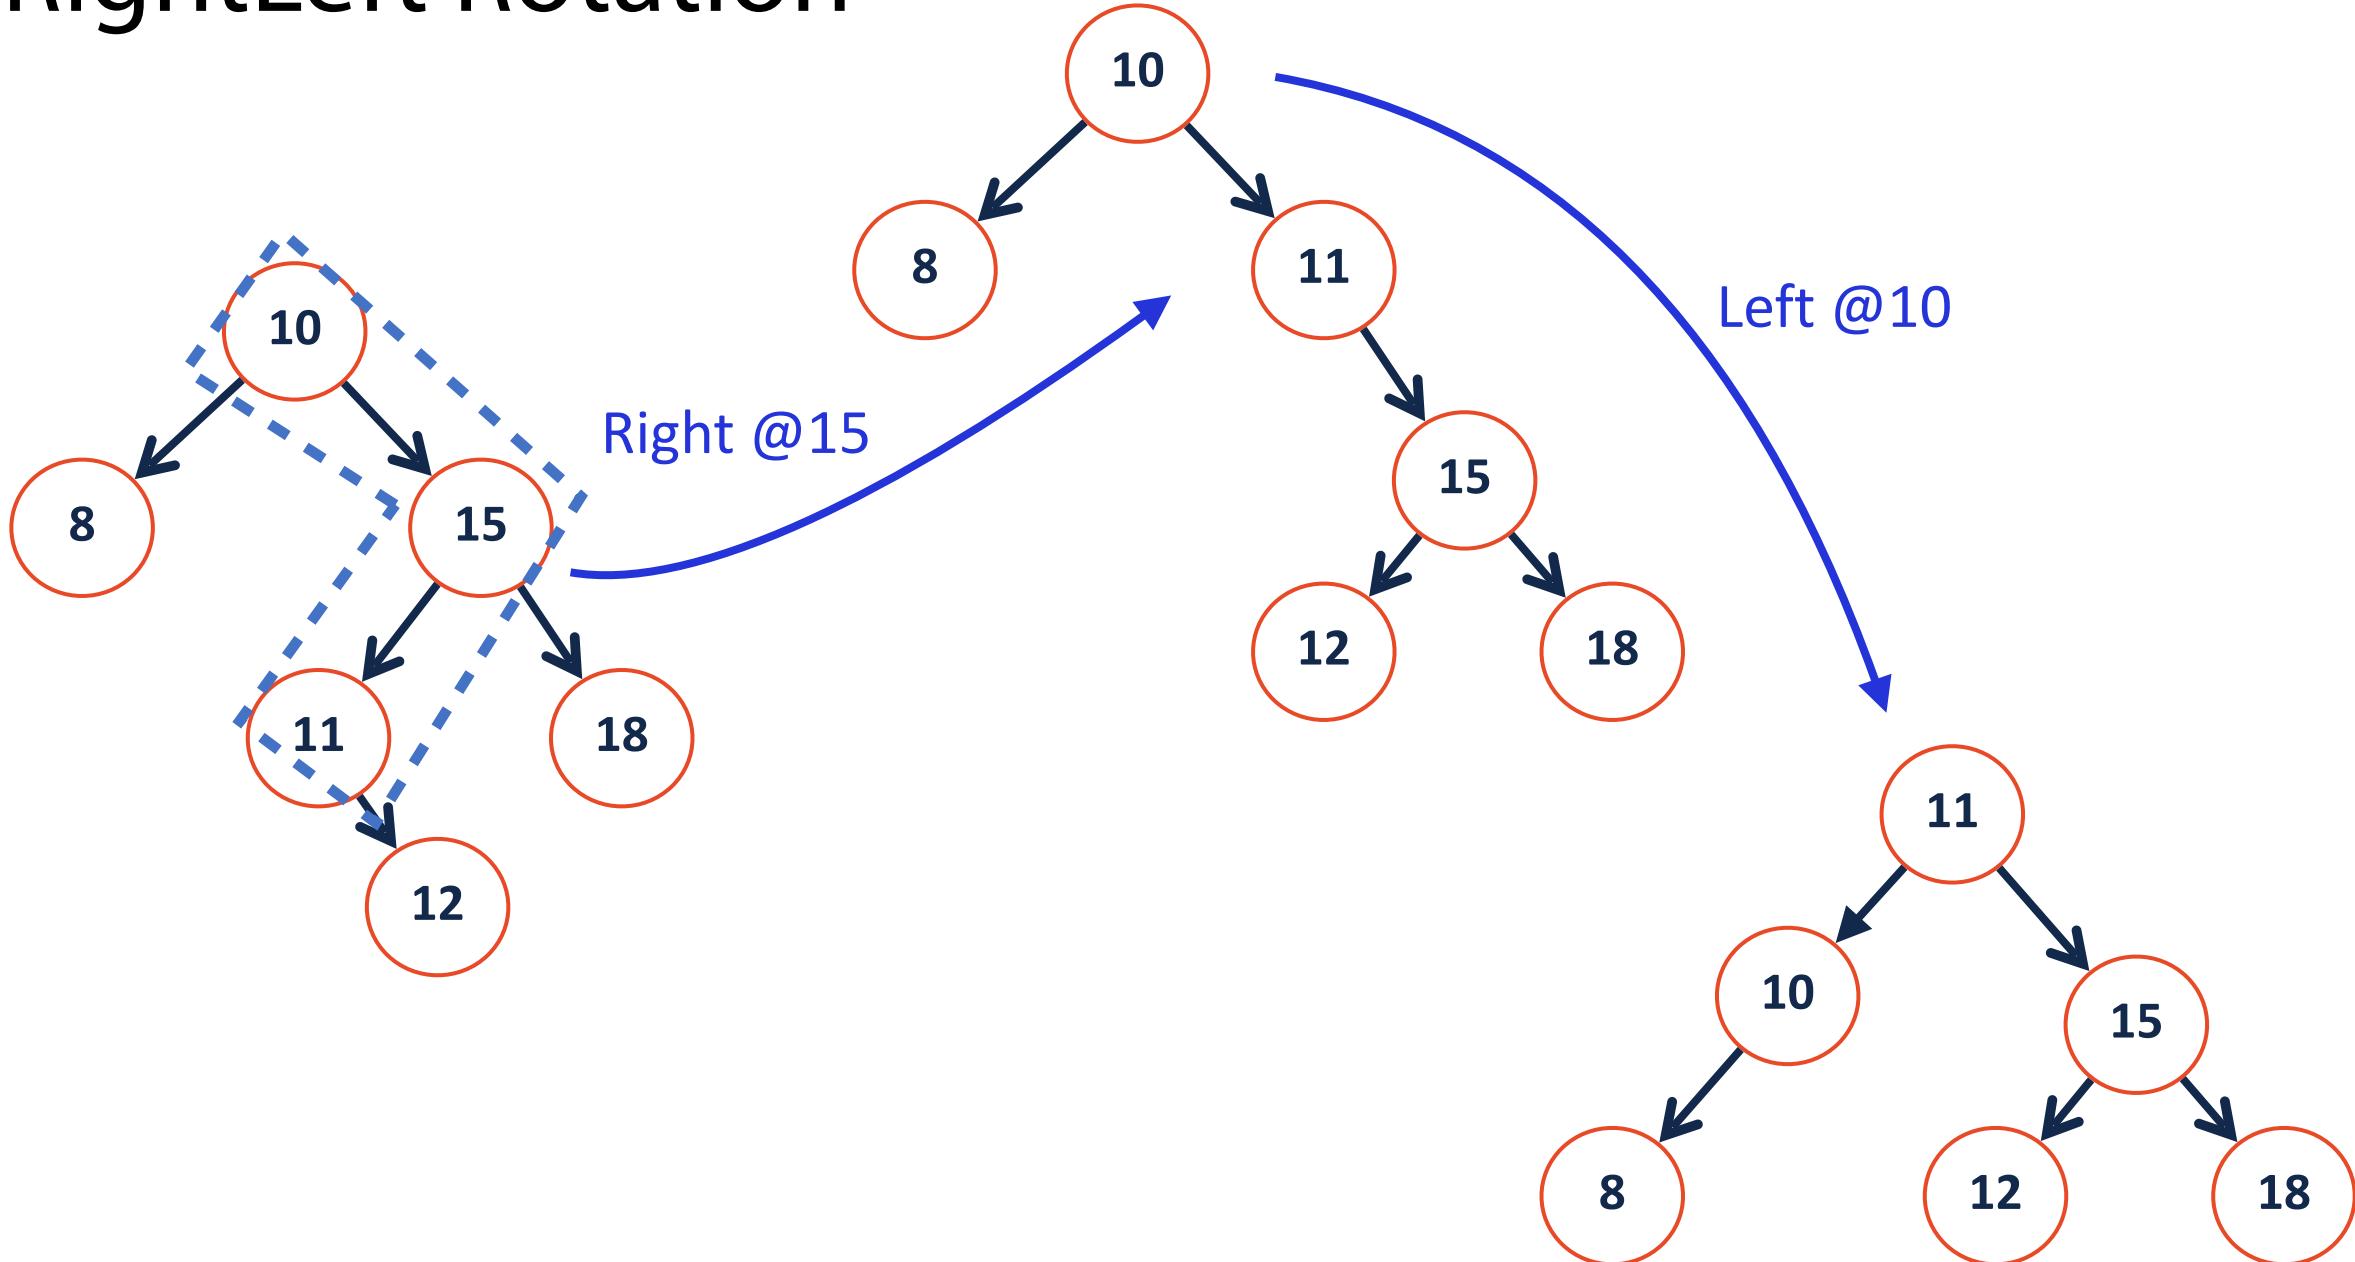
- 1) Create a tmp pointer to root
- 2) Update root to point to mid
- 3) `tmp->left = root->right`
- 4) `root->right = tmp`



# LeftRight Rotation



# RightLeft Rotation



# AVL Rotations



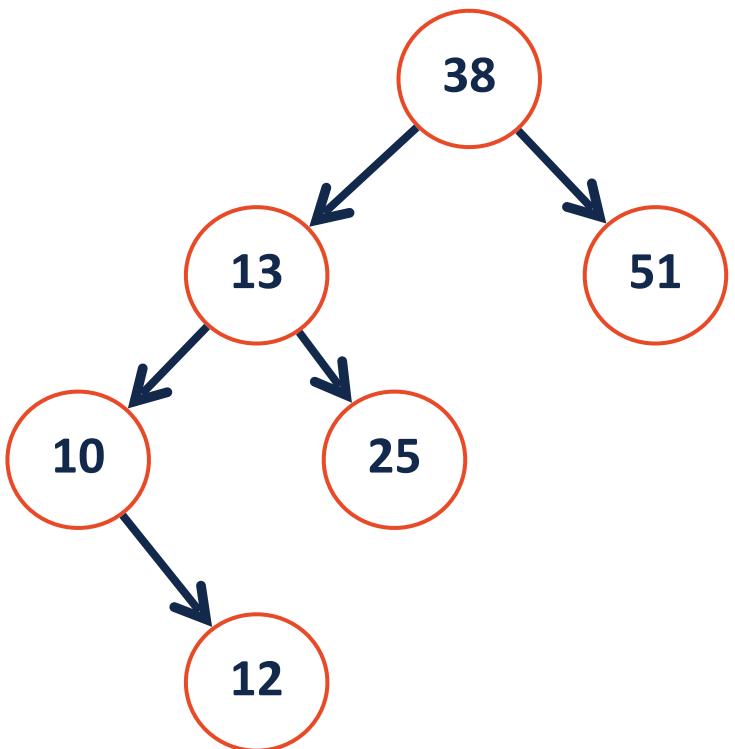
Four kinds of rotations: (L, R, LR, RL)

1. All rotations are local (subtrees are not impacted)
2. The running time of rotations are constant
3. The rotations maintain BST property

**Goal:**

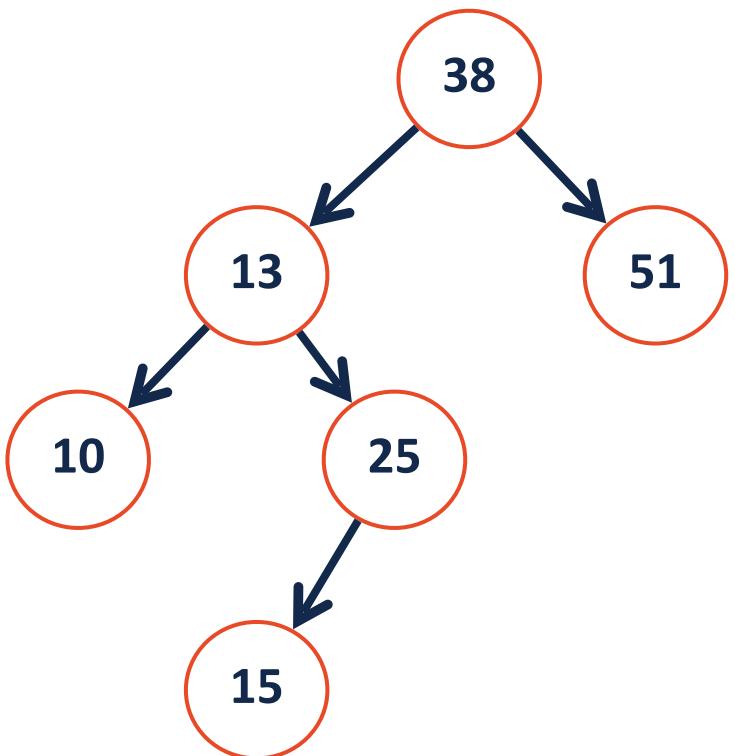
# AVL Rotations

We can identify which rotation to do using **balance**



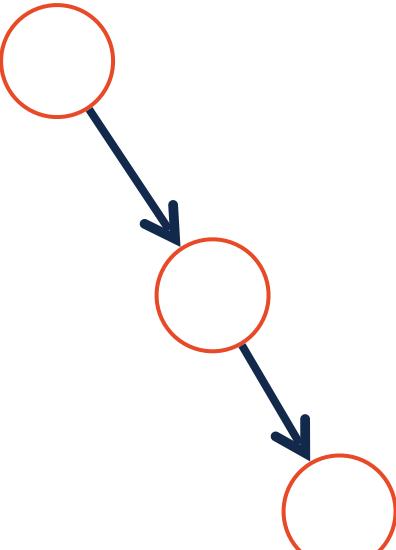
# AVL Rotations

We can identify which rotation to do using **balance**

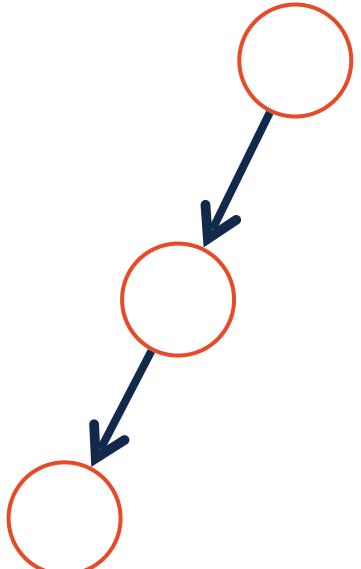


# AVL Rotations

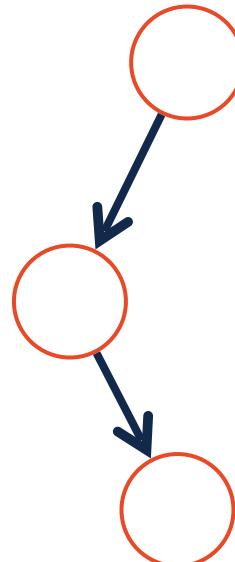
**Left**



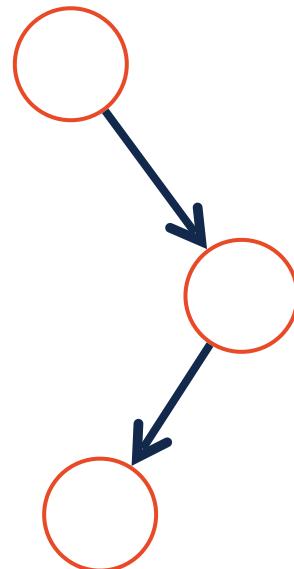
**Right**



**LeftRight**



**RightLeft**



Root Balance: 2

-2

-2

2

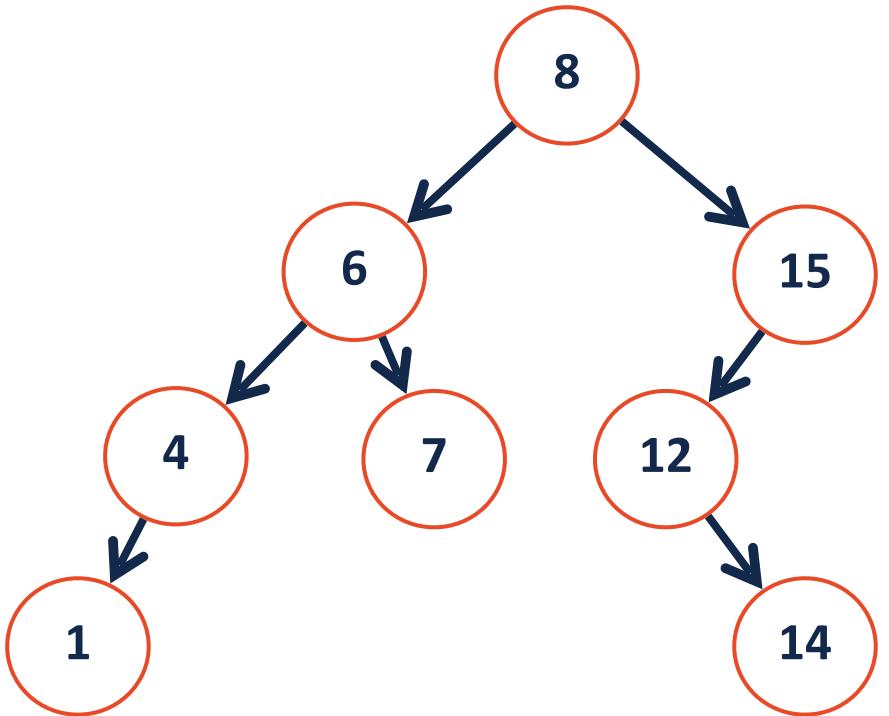
Child Balance: 1

-1

1

-1

# AVL Rotation Practice



# AVL vs BST ADT



The AVL tree is a modified binary search tree that rotates **when necessary**

```
1 struct TreeNode {  
2     T key;  
3     unsigned height;  
4     TreeNode *left;  
5     TreeNode *right;  
6 }
```

How does the constraint on balance affect the core functions?

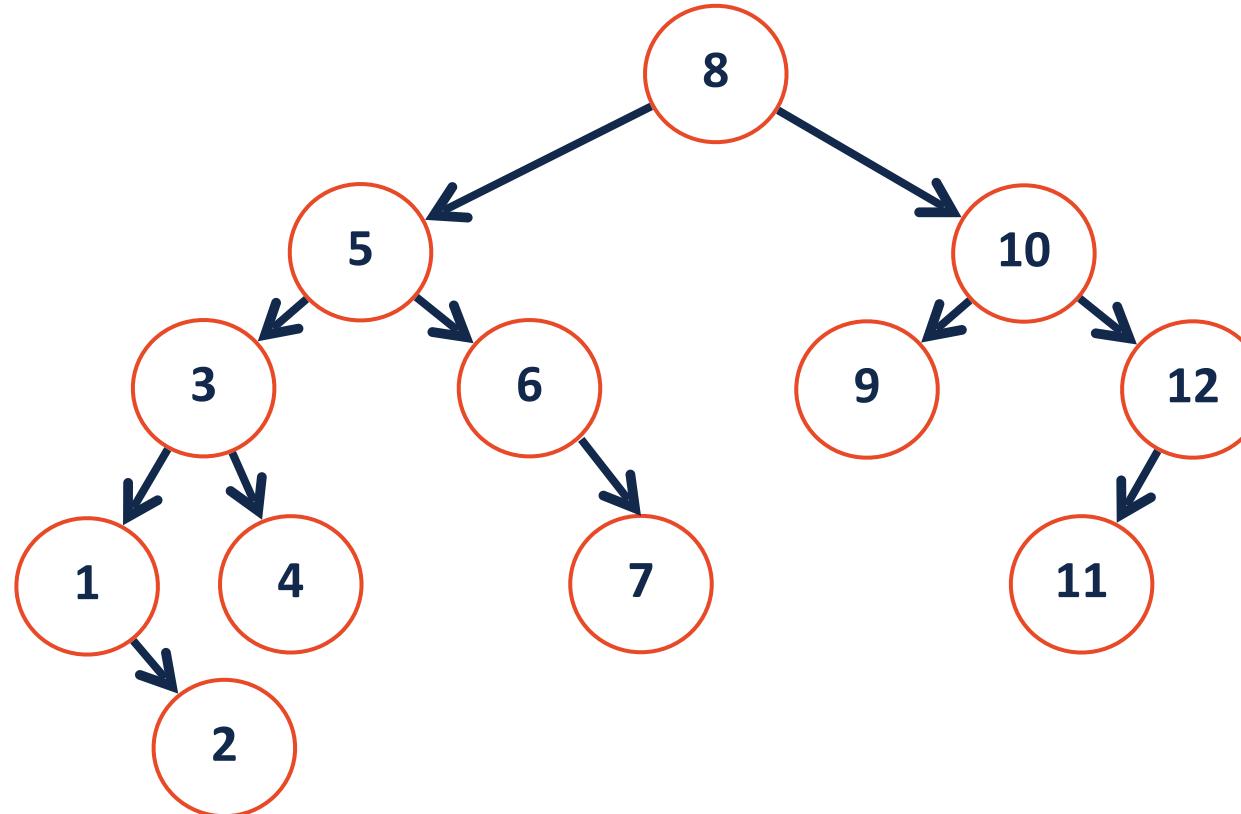
**Find**

**Insert**

**Remove**

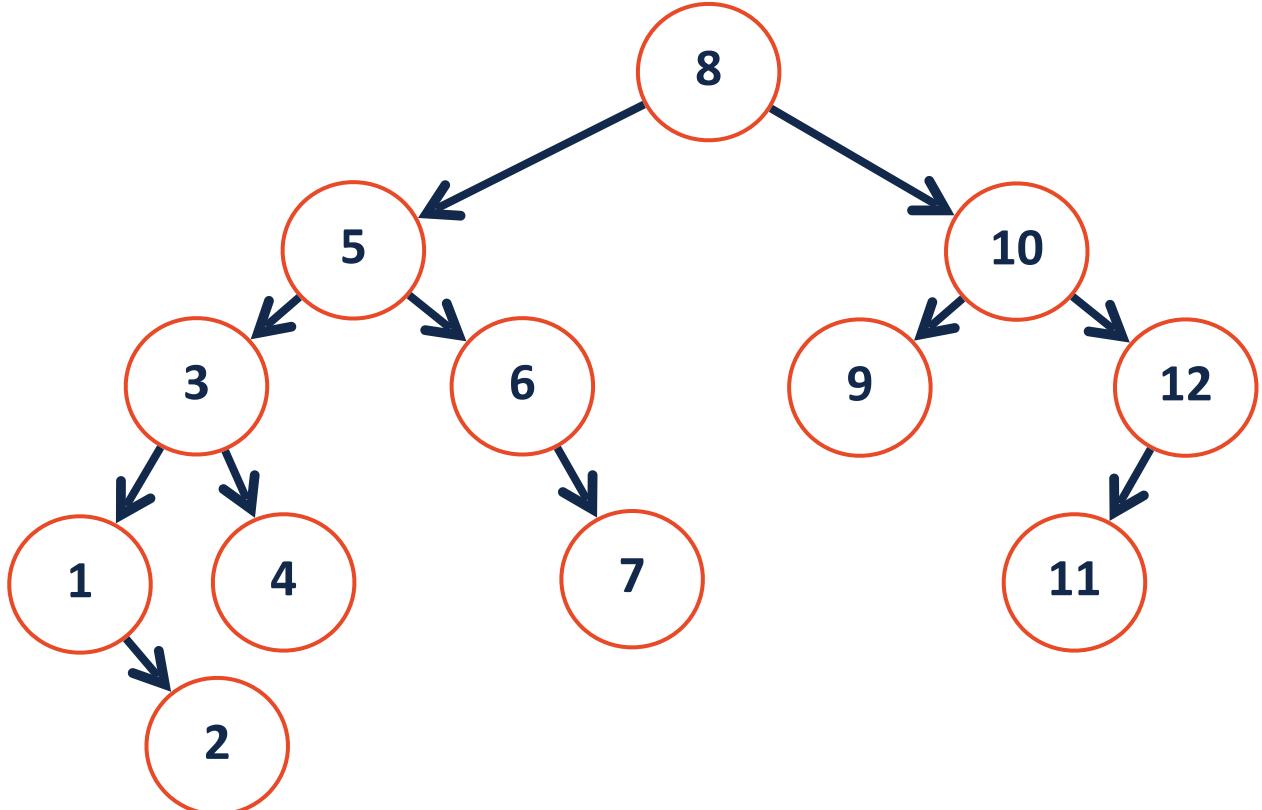
# AVL Find

\_find(7)



# AVL Insertion

\_insert(6.5)



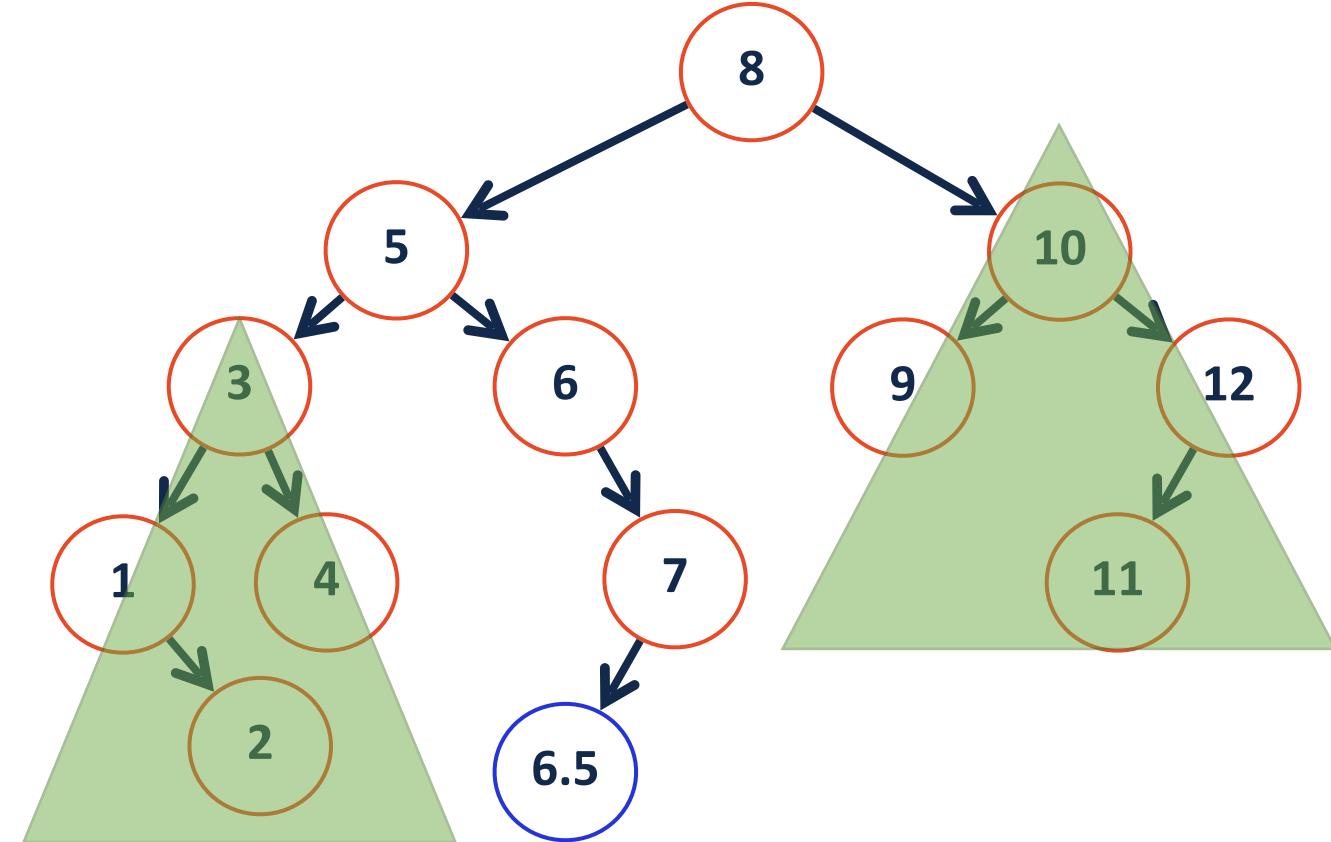
```
1 struct TreeNode {  
2     T key;  
3     unsigned height;  
4     TreeNode *left;  
5     TreeNode *right;  
6 };
```

# AVL Insertion

—insert(6.5)

## Insert (recursive pseudocode):

1. Insert at proper place
2. Check for imbalance
3. Rotate, if necessary
4. Update height



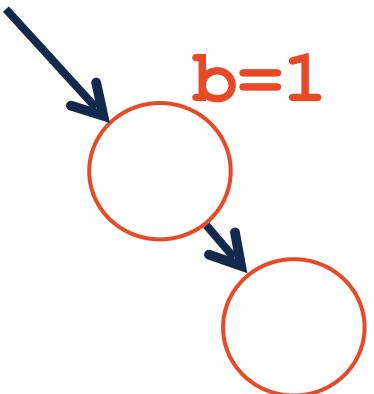
```
1 struct TreeNode {  
2     T key;  
3     unsigned height;  
4     TreeNode *left;  
5     TreeNode *right;  
6 };
```

```
151 template <typename K, typename V>
152 void AVL<K, D>::_insert(const K & key, const V & data, TreeNode
*& cur) {
153     if (cur == NULL)           { cur = new TreeNode(key, data);    }
157     else if (key < cur->key) { _insert( key, data, cur->left ); }
160     else if (key > cur->key) { _insert( key, data, cur->right ); }
166     _ensureBalance(cur);
167 }
```

```
119 template <typename K, typename V>
120 void AVL<K, D>::_ensureBalance(TreeNode *& cur) {
121     // Calculate the balance factor:
122     int balance = height(cur->right) - height(cur->left);
123
124     // Check if the node is current not in balance:
125     if ( balance == -2 ) {
126         int l_balance =
127             height(cur->left->right) - height(cur->left->left);
128         if ( l_balance == -1 ) { _____; }
129         else { _____; }
130     } else if ( balance == 2 ) {
131         int r_balance =
132             height(cur->right->right) - height(cur->right->left);
133         if( r_balance == 1 ) { _____; }
134         else { _____; }
135     }
136     _updateHeight(cur);
137 }
```

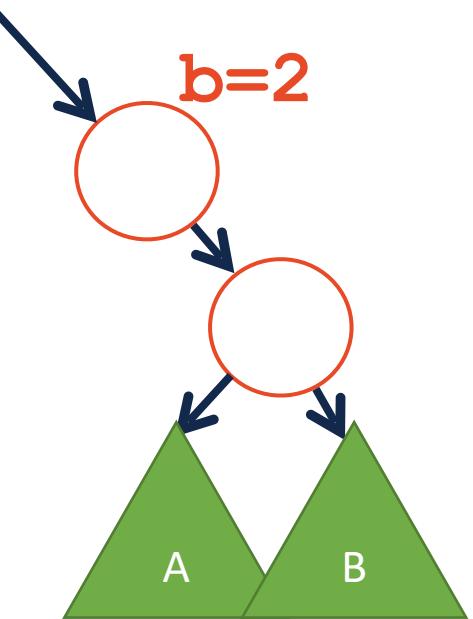
# AVL Insertion

Given an AVL is balanced, insert can create **at most** one imbalance



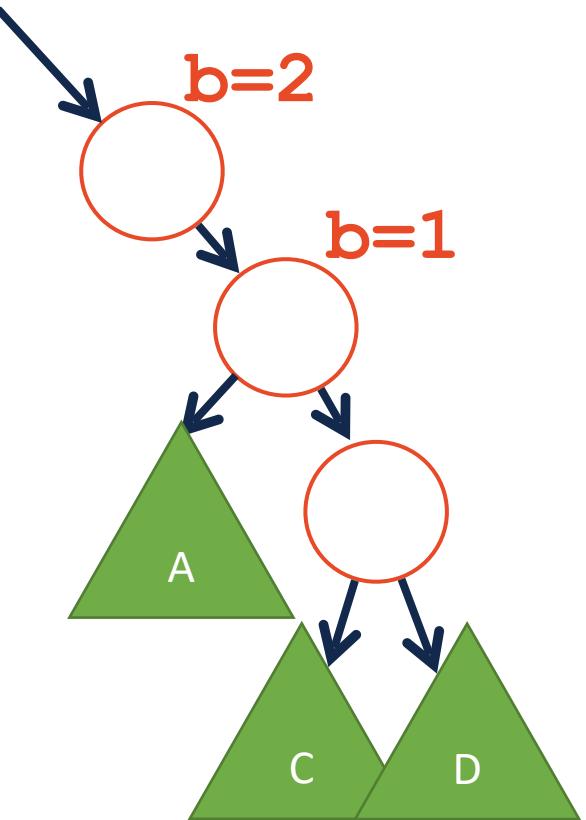
# AVL Insertion

Given an AVL is balanced, insert can create **at most** one imbalance



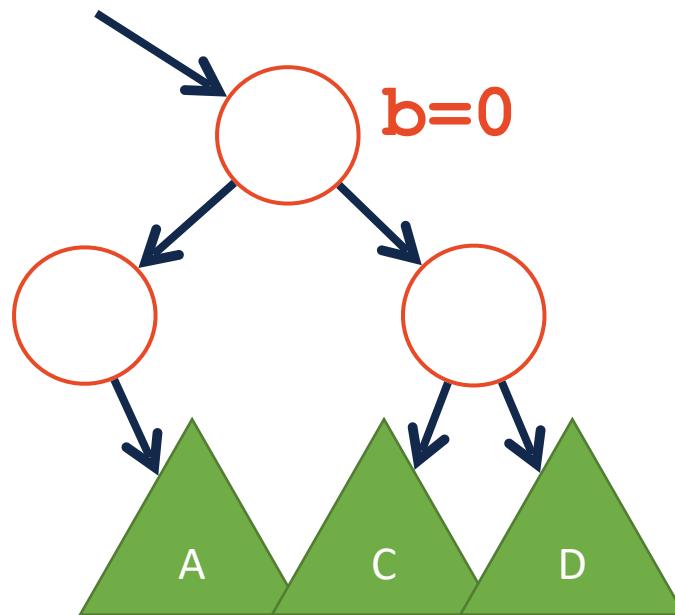
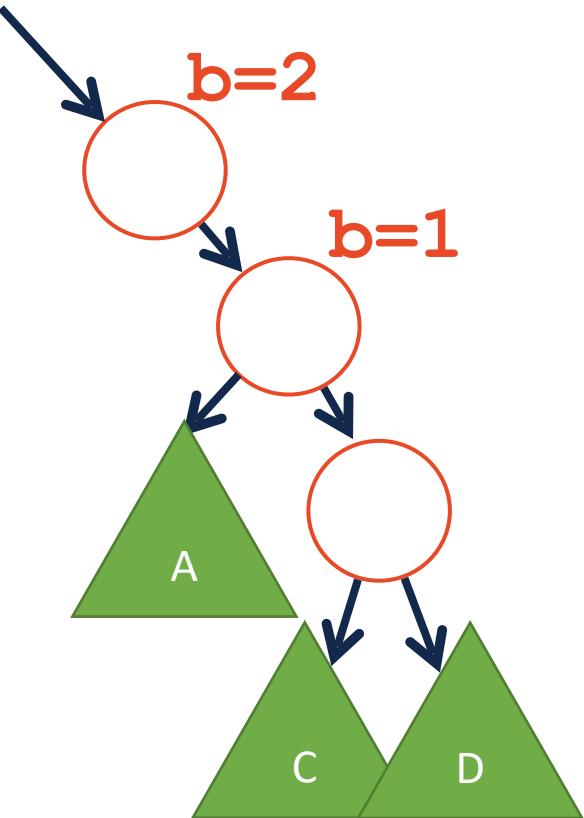
# AVL Insertion

If we insert in B, I must have a balance pattern of 2, 1



# AVL Insertion

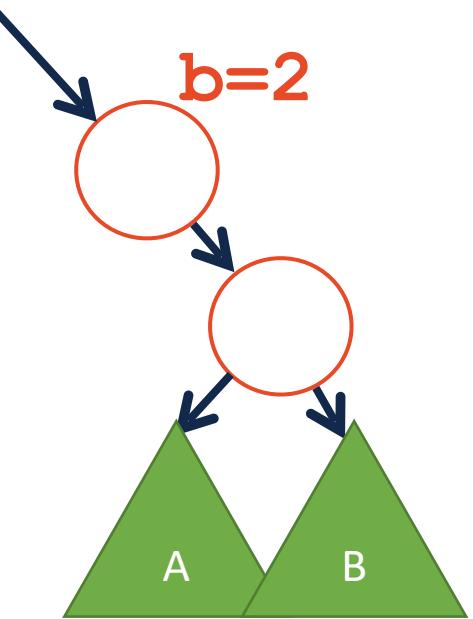
A **left** rotation fixes our imbalance in our local tree.



After rotation, subtree has **pre-insert height**. (Overall tree is balanced)

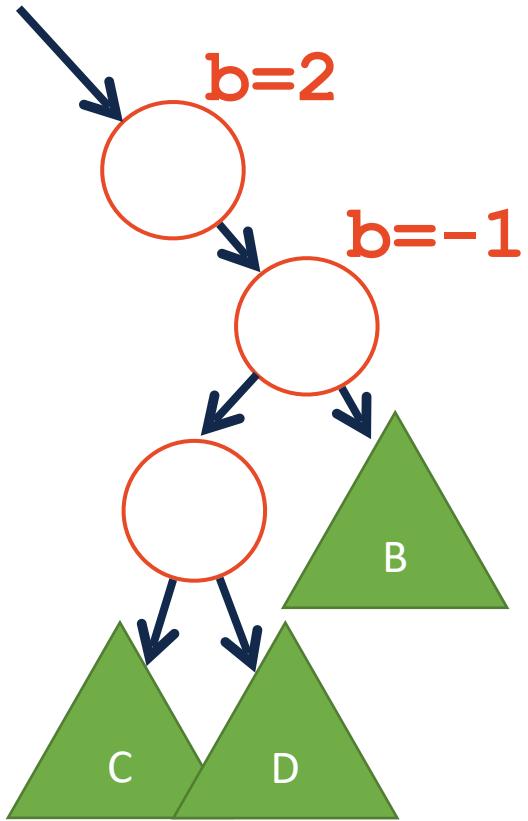
# AVL Insertion

If we insert in A, I must have a balance pattern of **2, -1**



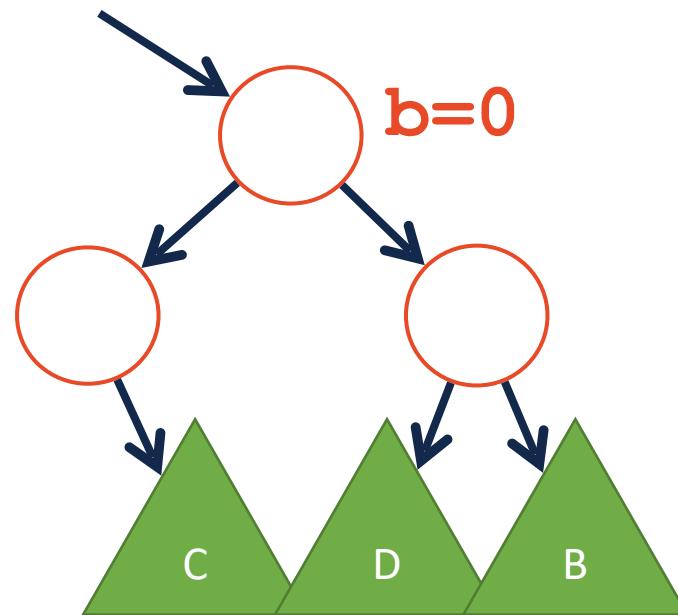
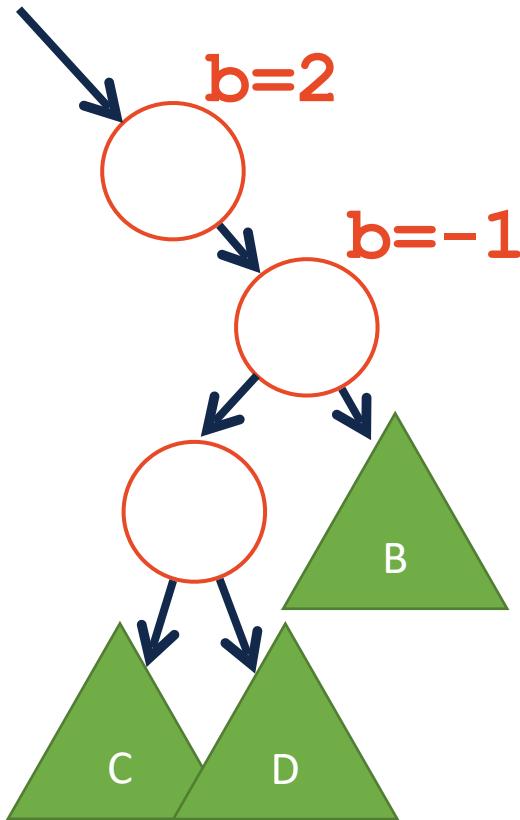
# AVL Insertion

If we insert in A, I must have a balance pattern of 2, -1



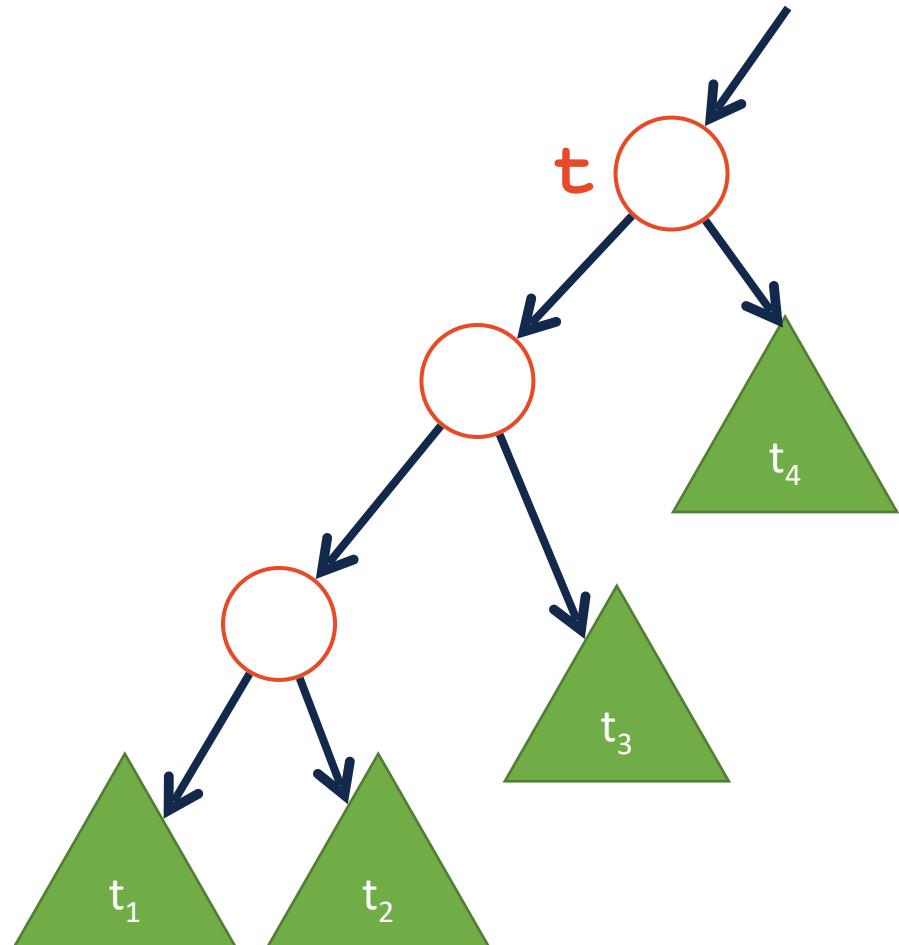
# AVL Insertion

A **rightLeft** rotation fixes our imbalance in our local tree.



After rotation, subtree has **pre-insert height**. (Overall tree is balanced)

# AVL Insertion

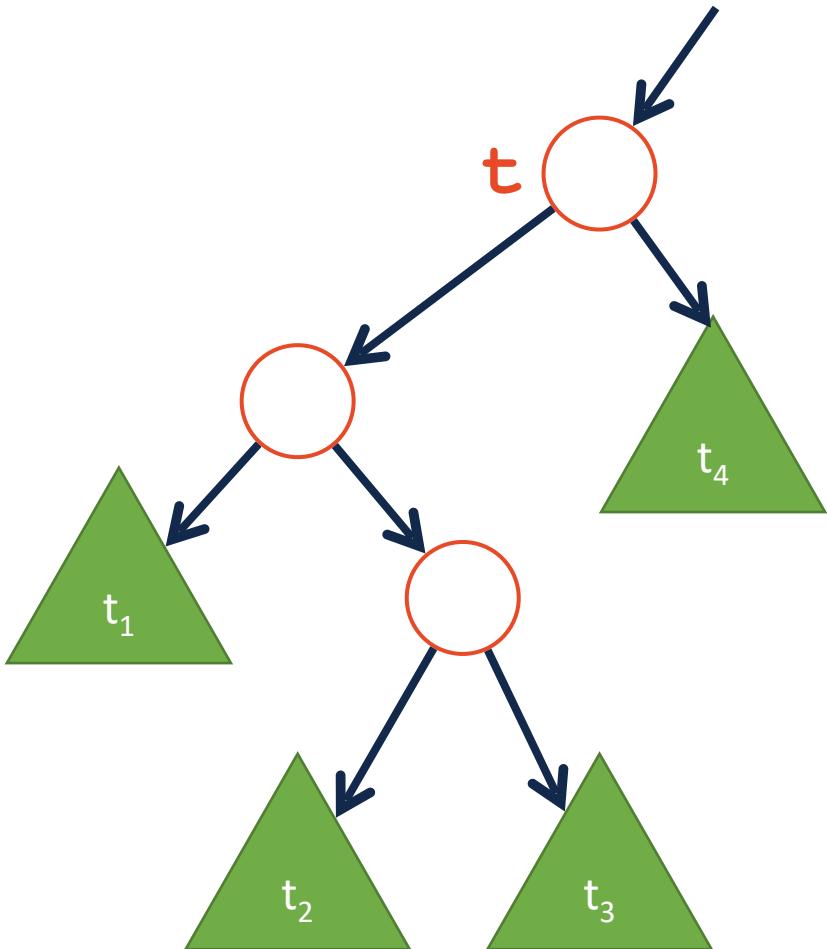


## Theorem:

If an insertion occurred in subtrees  $t_1$  or  $t_2$  and an imbalance was first detected at  $t$ , then a \_\_\_\_\_ rotation about  $t$  restores the balance of the tree.

We gauge this by noting the balance factor of  $t$  is \_\_\_\_\_ and the balance factor of  $t\text{-}left$  is \_\_\_\_\_.

# AVL Insertion



## Theorem:

If an insertion occurred in subtrees  $t_2$  or  $t_3$  and an imbalance was first detected at  $t$ , then a \_\_\_\_\_ rotation about  $t$  restores the balance of the tree.

We gauge this by noting the balance factor of  $t$  is \_\_\_\_\_ and the balance factor of  $t\text{-}left$  is \_\_\_\_\_.

# AVL Insertion



We've seen every possible insert that can cause an imbalance

Insert *may* increase height by at most:

A rotation reduces the height of the subtree by:

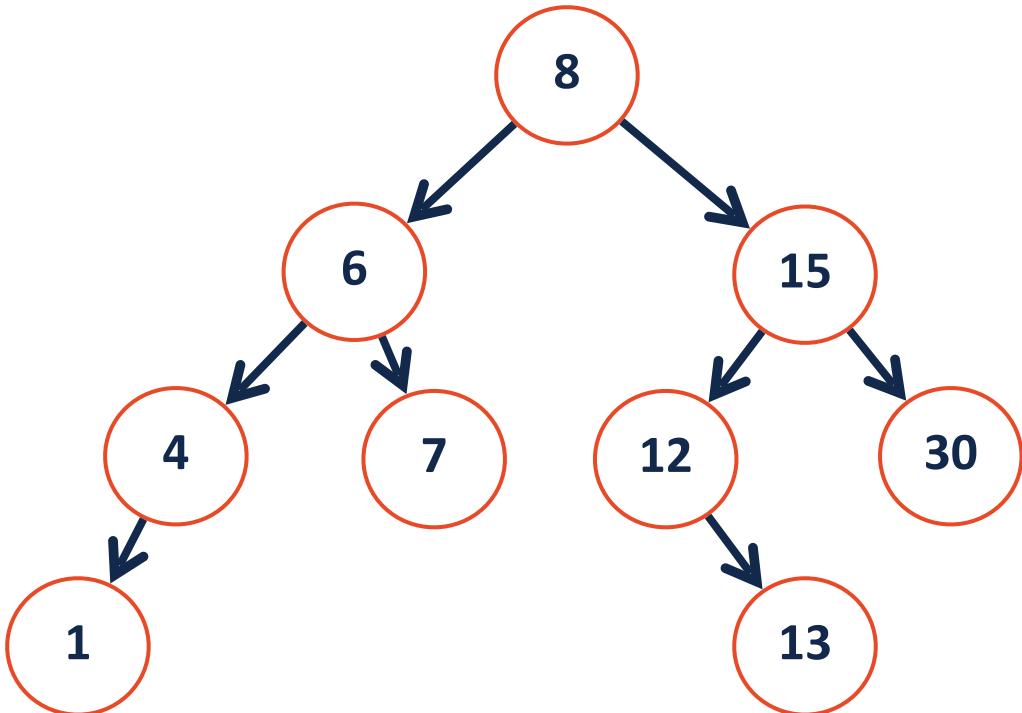
**A single\* rotation restores balance and corrects height!**

What is the Big O of performing our rotation?

What is the Big O of insert?

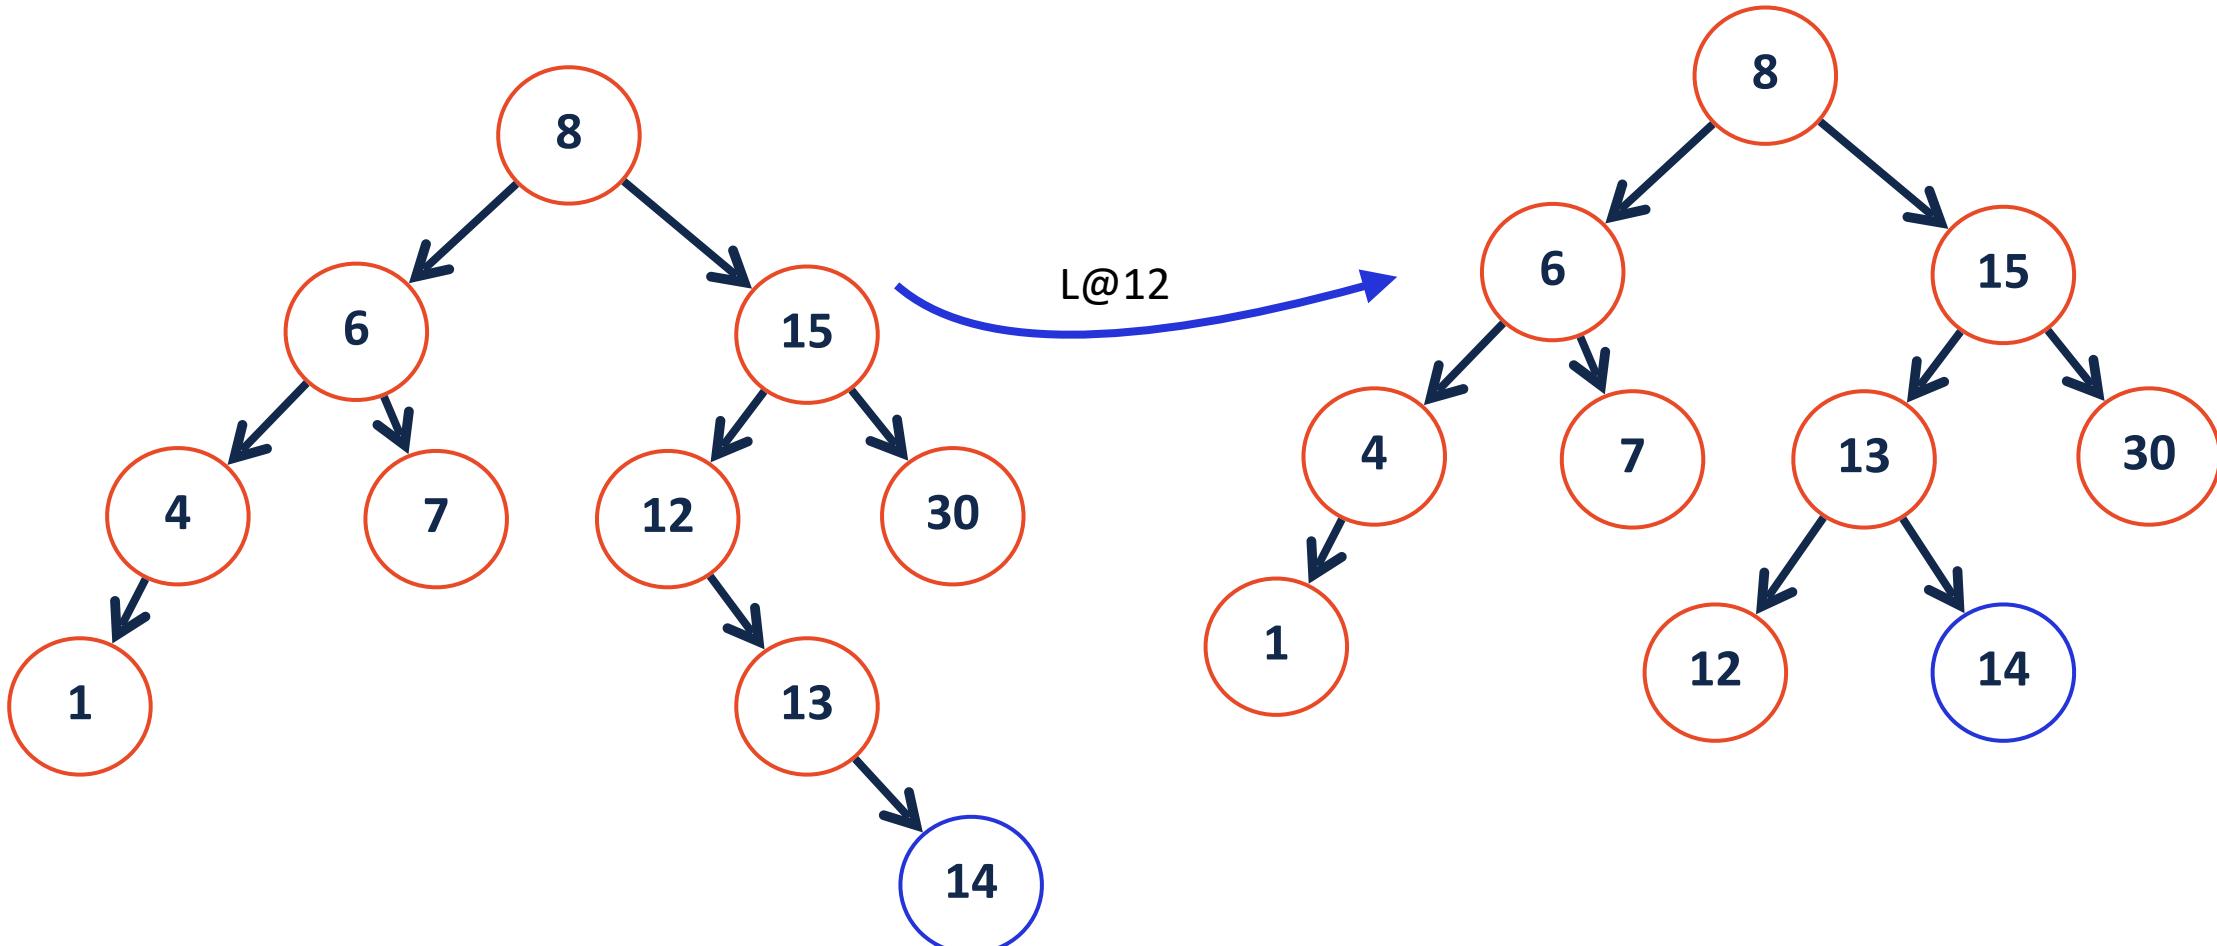
# AVL Insertion Practice

\_insert(14)



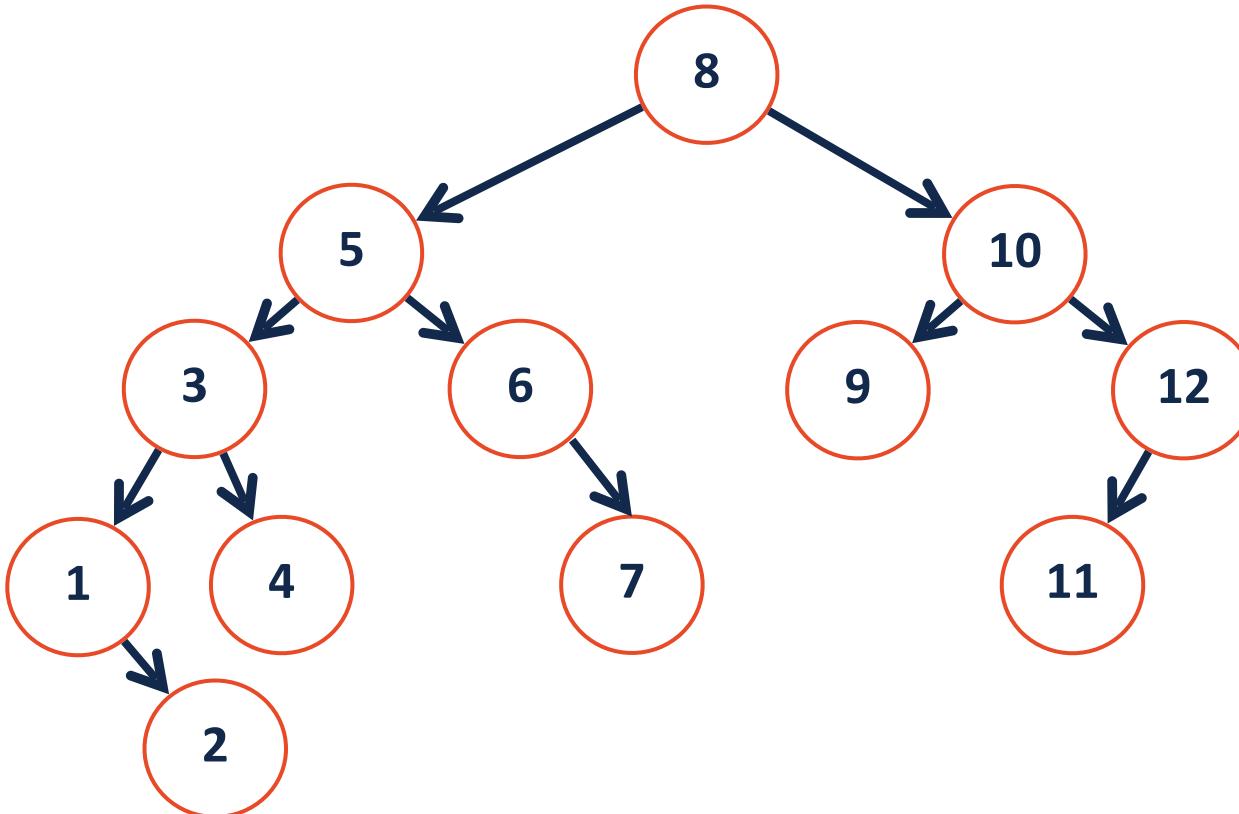
# AVL Insertion Practice

\_insert(14)



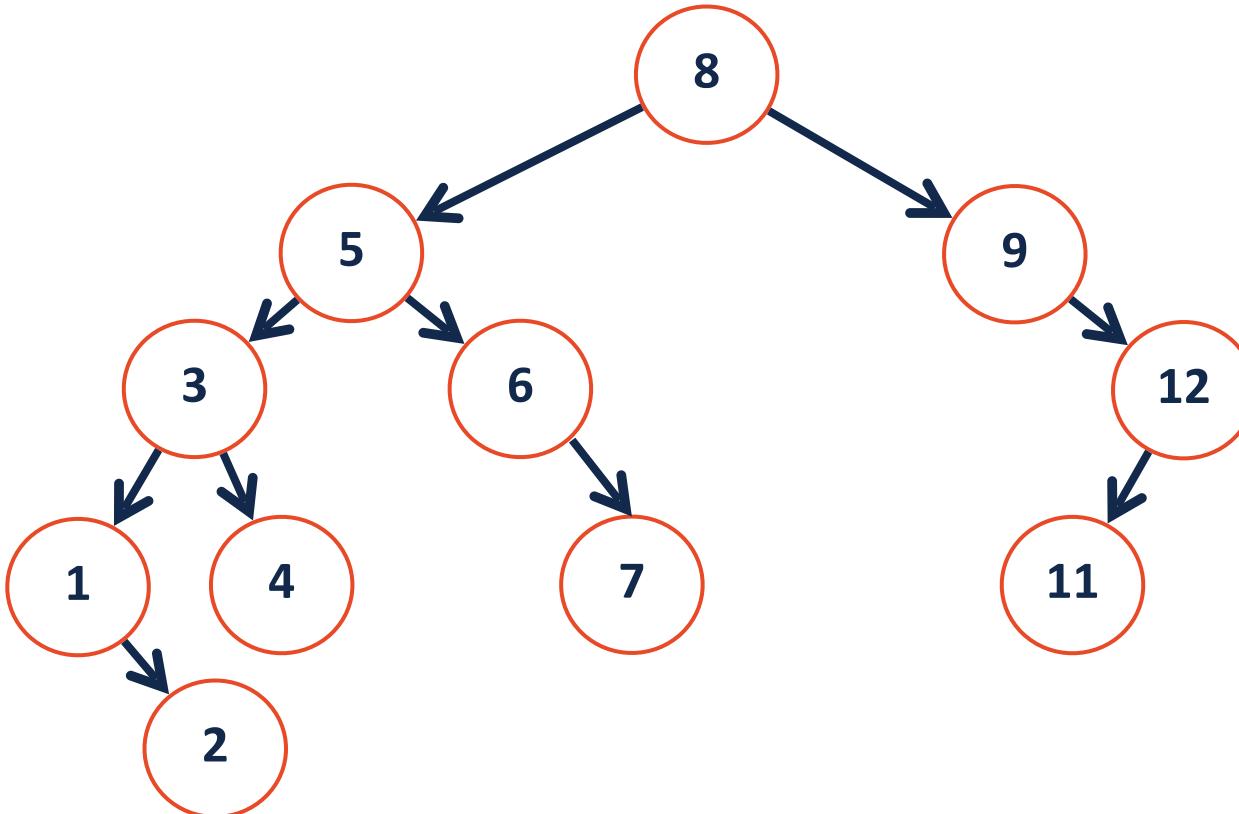
# AVL Remove

\_remove(10)



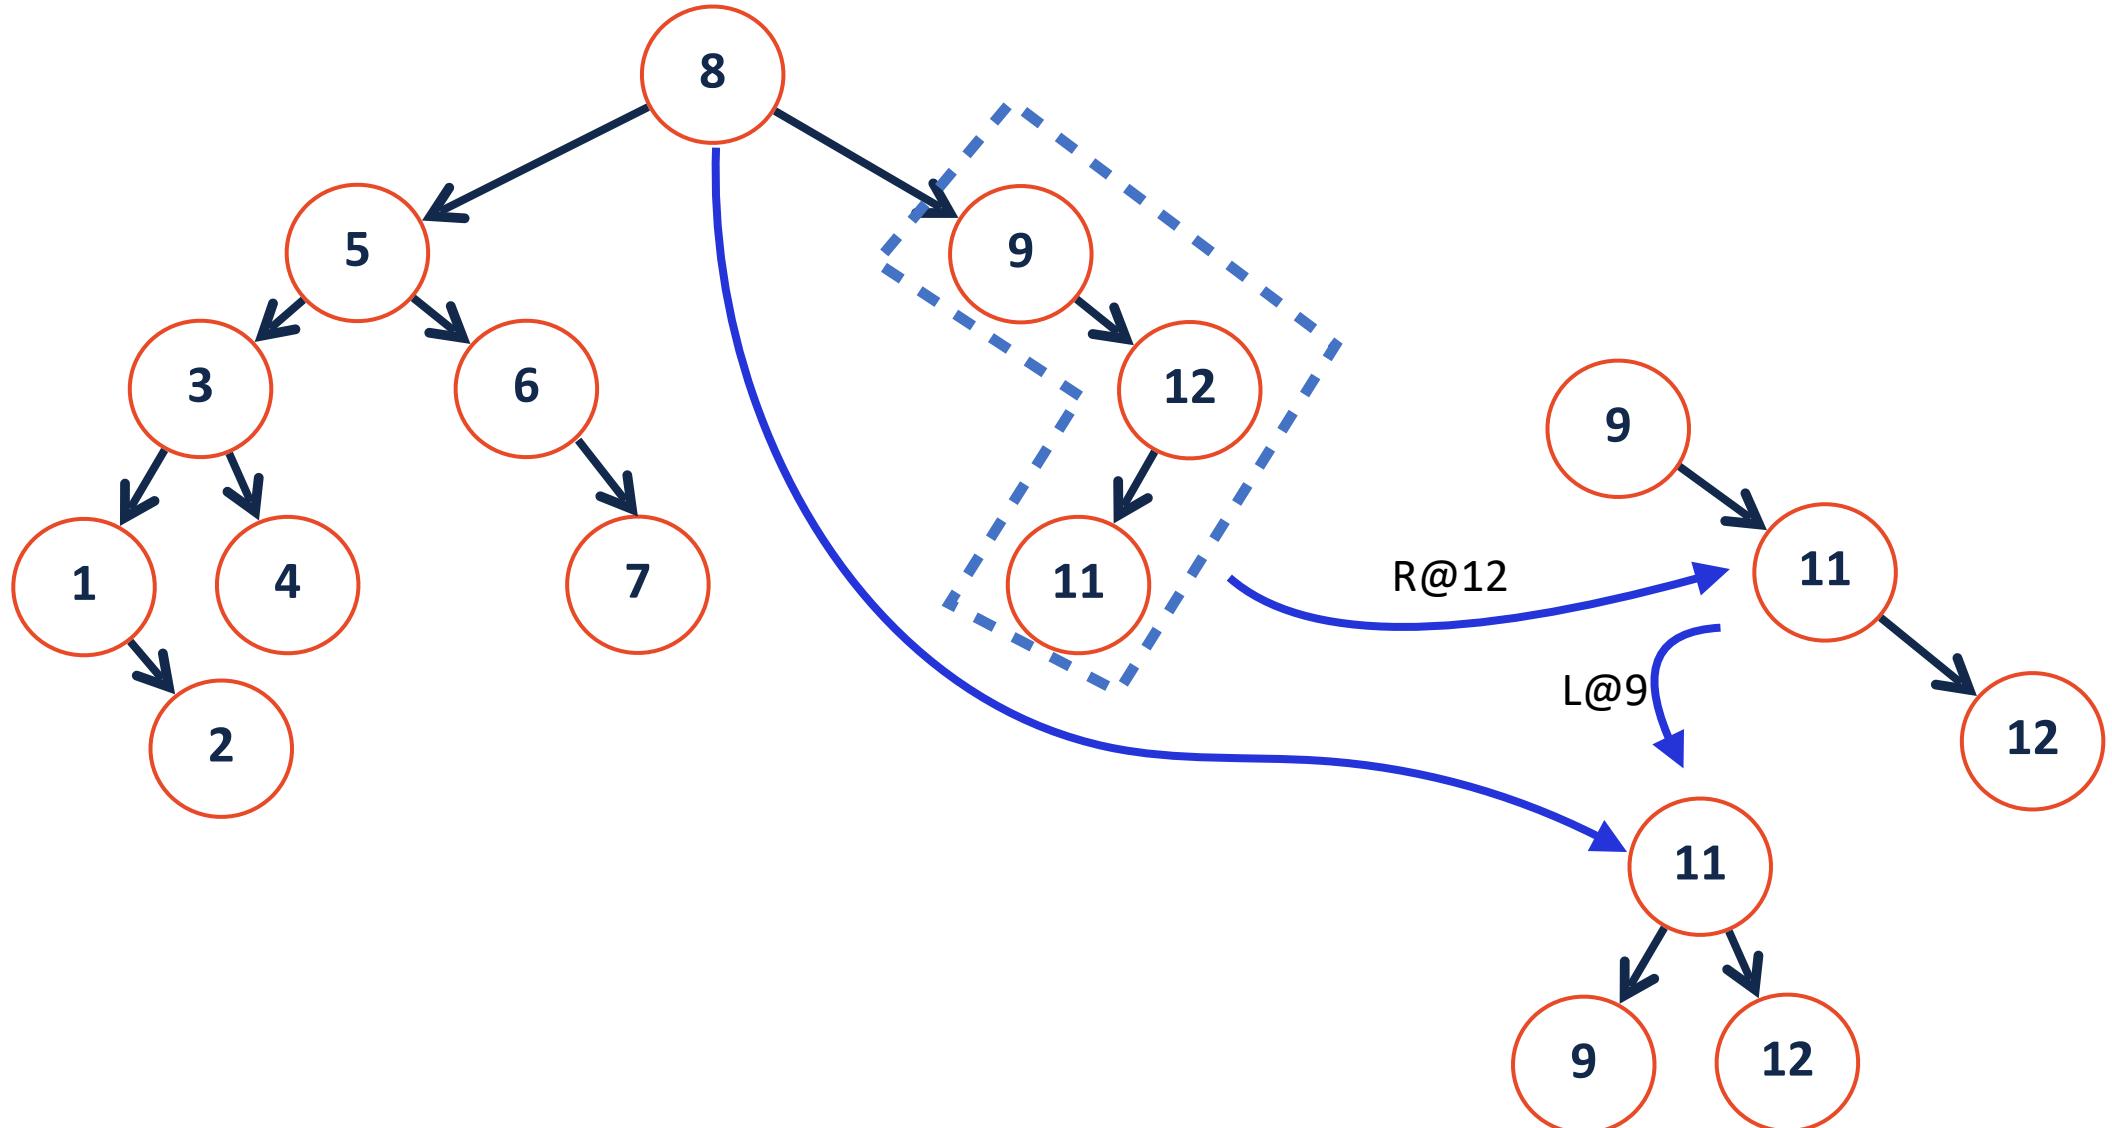
# AVL Remove

`_remove(10)`



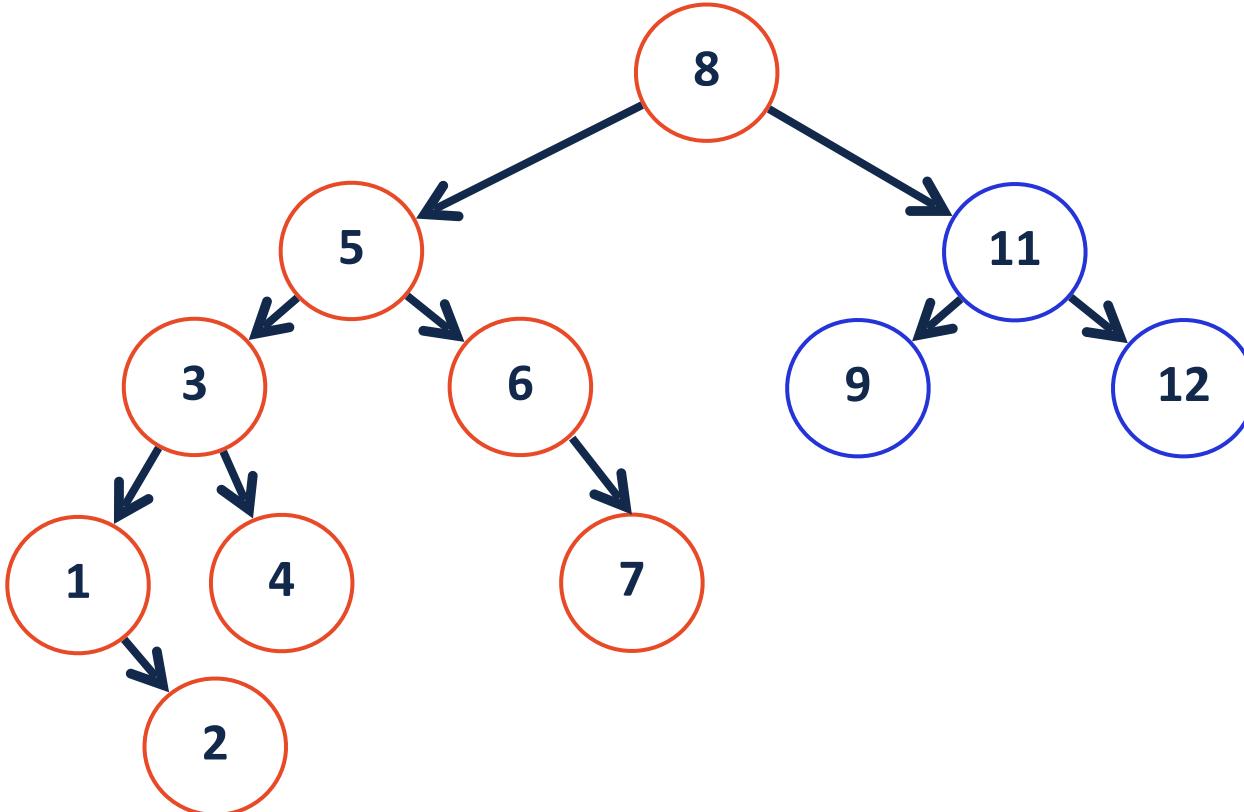
# AVL Remove

\_remove(10)



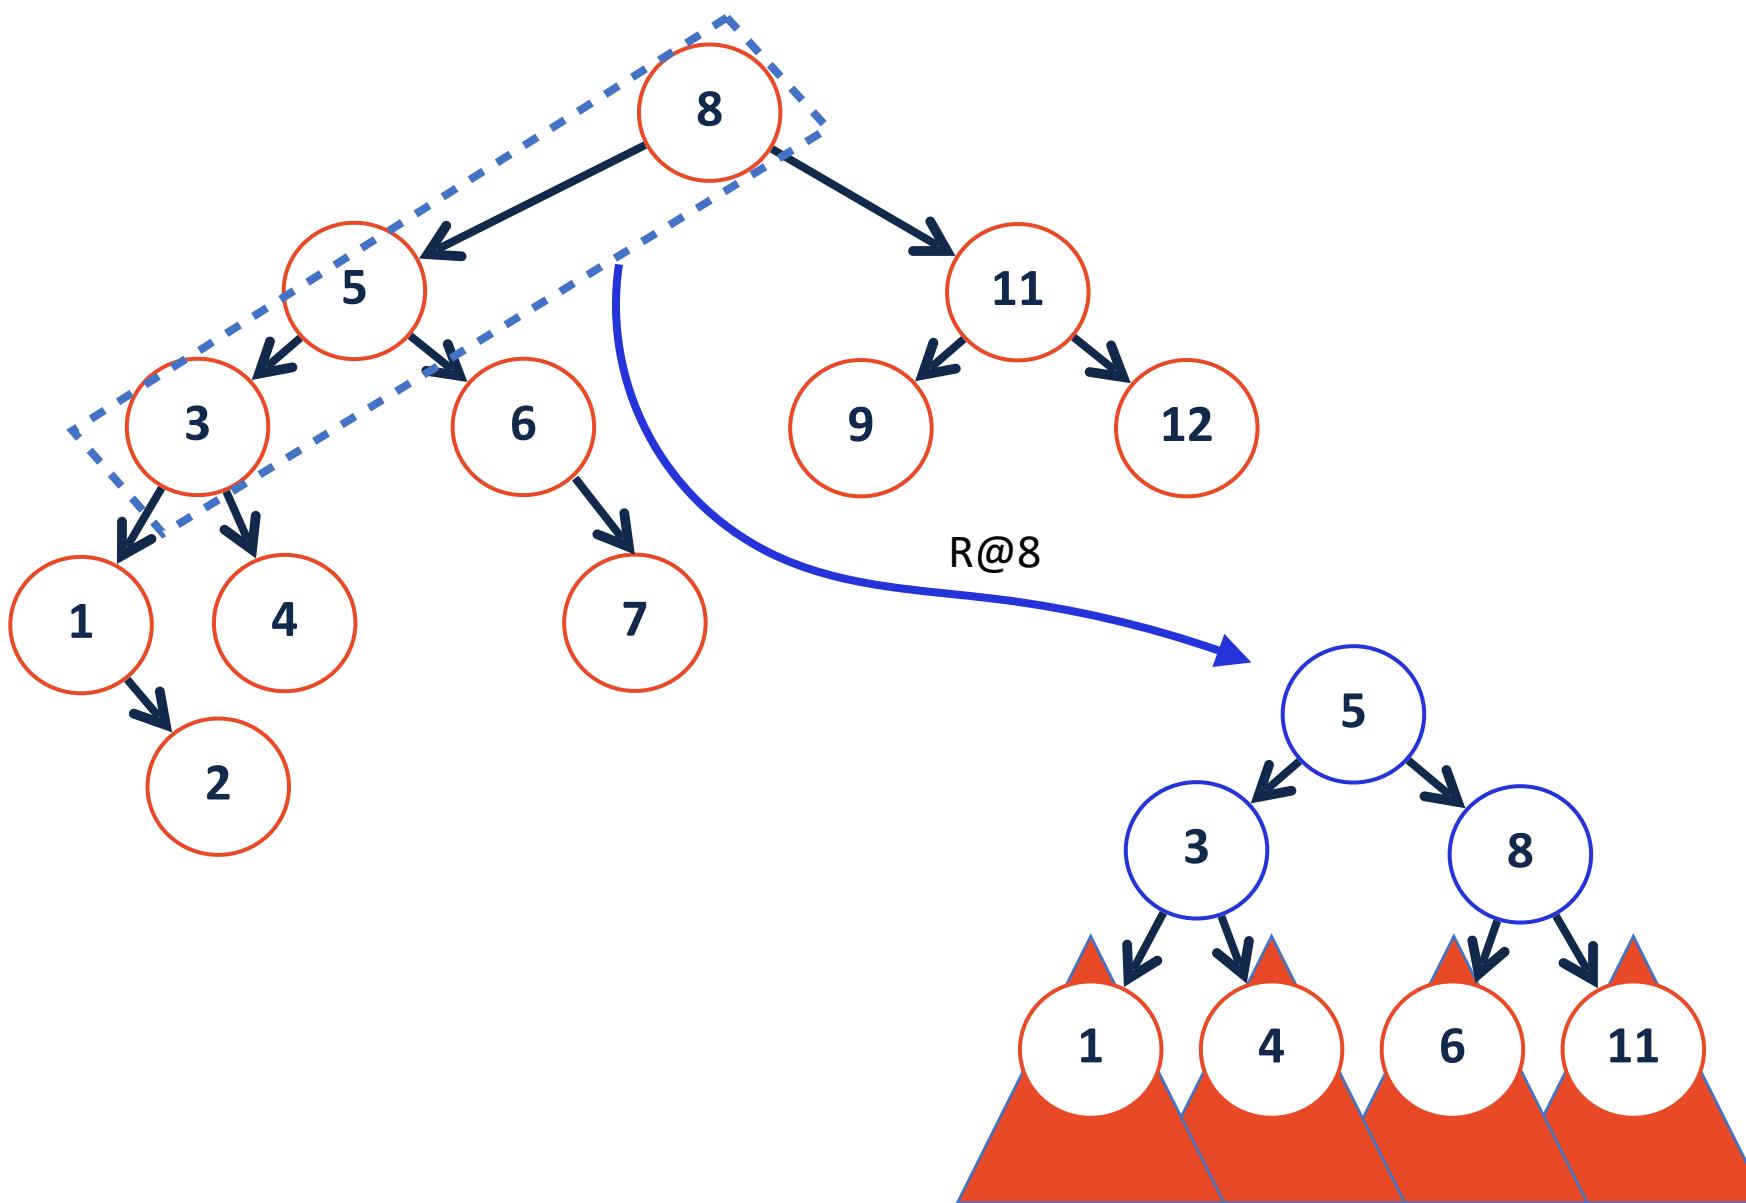
# AVL Remove

\_remove(10)



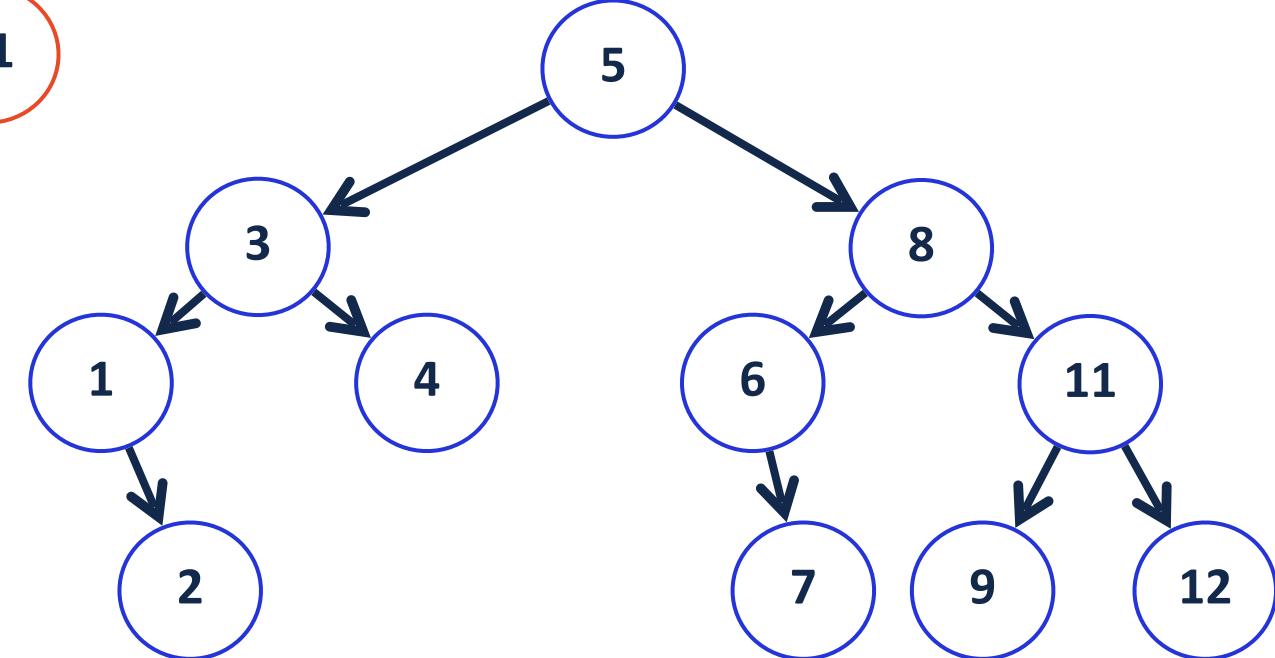
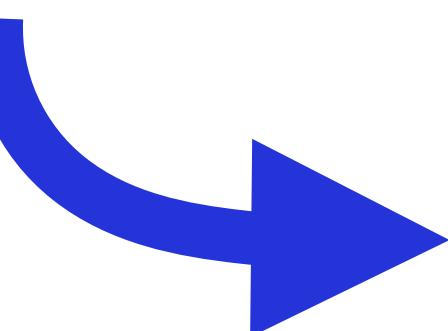
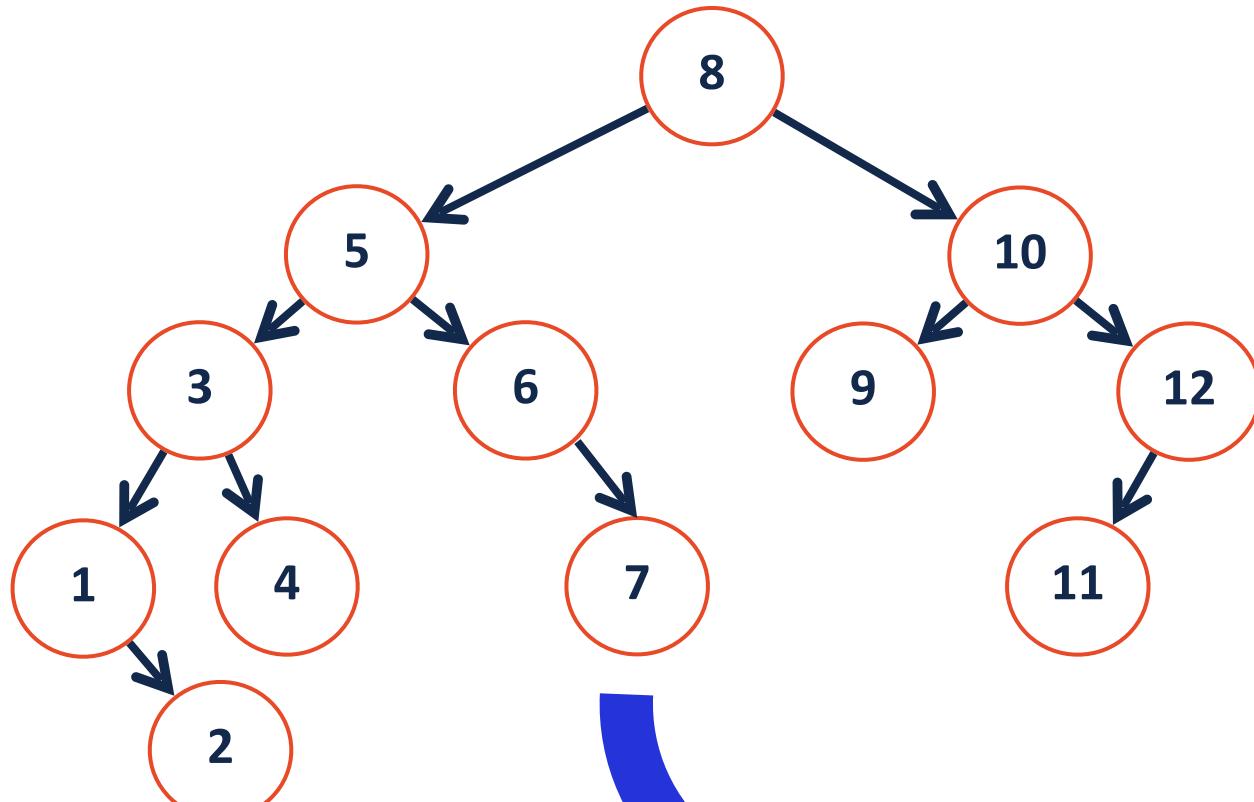
# AVL Remove

\_remove(10)



# AVL Remove

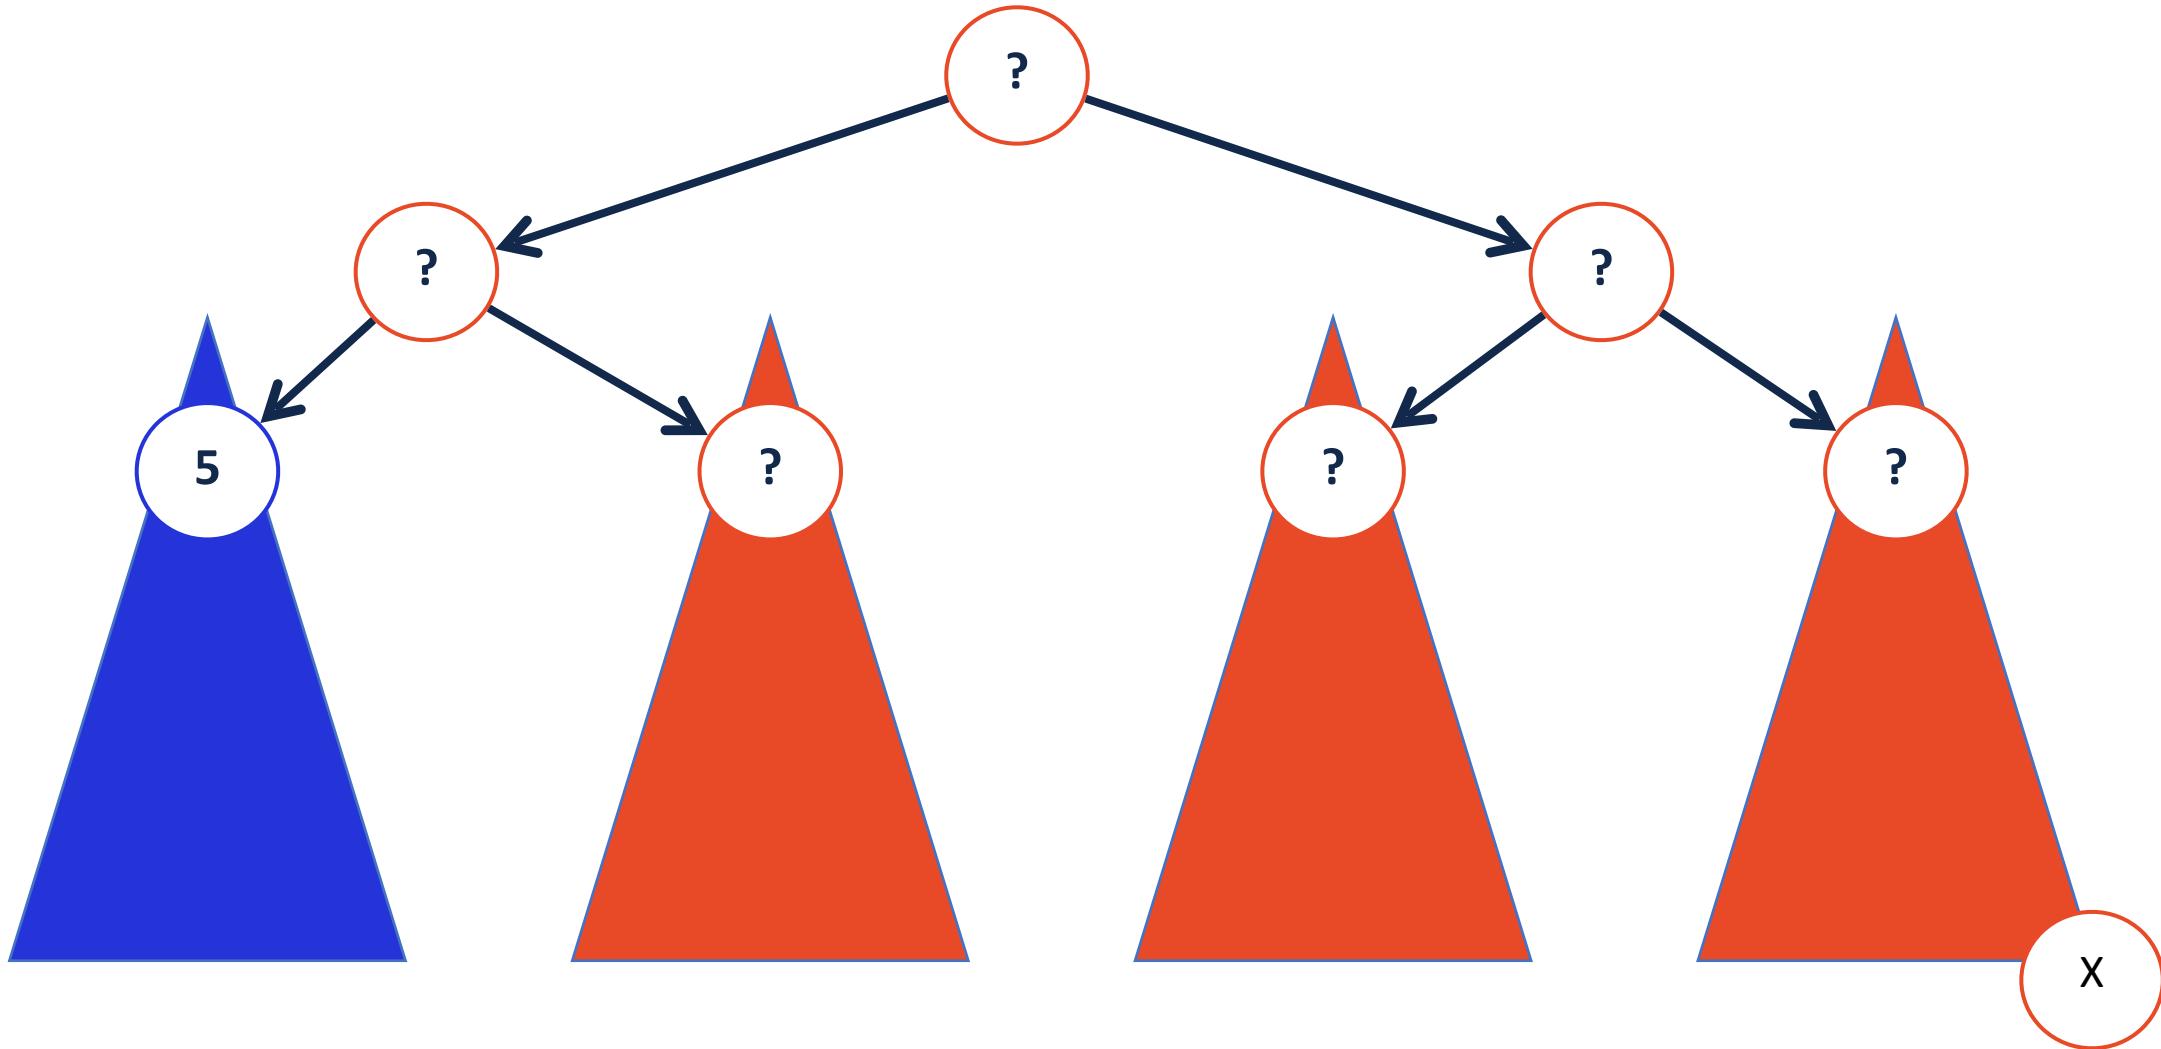
remove (10) ⏱



**Remove (pseudo code):**

- 1: Remove at proper place
- 2: Check for imbalance
- 3: Rotate, if necessary
- 4: Update height

# AVL Remove



# AVL Remove



An AVL remove step can reduce a subtree height by at most:

But a rotation *reduces* the height of a subtree by one!

**We might have to perform a rotation at every level of the tree!**

# AVL Tree Analysis

For an AVL tree of height  $h$ :

Find runs in: \_\_\_\_\_.

Insert runs in: \_\_\_\_\_.

Remove runs in: \_\_\_\_\_.

**Claim:** The height of the AVL tree with  $n$  nodes is: \_\_\_\_\_.